

Convergence Properties of Kronecker Graphical Lasso Algorithms

Theodoros Tsiligkaridis *, *Student Member, IEEE*, Alfred O. Hero III, *Fellow, IEEE*, Shuheng Zhou, *Member, IEEE*

Abstract

This report presents a thorough convergence analysis of Kronecker graphical lasso (KGLasso) algorithms for estimating the covariance of an i.i.d. Gaussian random sample under a sparse Kronecker-product covariance model. The KGLasso model, originally called the transposable regularized covariance model by Allen *et al* [1], implements a pair of ℓ_1 penalties on each Kronecker factor to enforce sparsity in the covariance estimator. The KGLasso algorithm generalizes Glasso, introduced by Yuan and Lin [2] and Banerjee *et al* [3], to estimate covariances having Kronecker product form. It also generalizes the unpenalized ML flip-flop (FF) algorithm of Dutilleul [4] and Werner *et al* [5] to estimation of sparse Kronecker factors. We establish that the KGLasso iterates converge pointwise to a local maximum of the penalized likelihood function. We derive high dimensional rates of convergence to the true covariance as both the number of samples and the number of variables go to infinity. Our results establish that KGLasso has significantly faster asymptotic convergence than FF and Glasso. Our results establish that KGLasso has significantly faster asymptotic convergence than FF and Glasso. Simulations are presented that validate the results of our analysis. For example, for a sparse $10,000 \times 10,000$ covariance matrix equal to the Kronecker product of two 100×100 matrices, the root mean squared error of the inverse covariance estimate using FF is 3.5 times larger than that obtainable using KGLasso.

Index Terms

Sparsity, structured covariance estimation, penalized maximum likelihood, graphical lasso, direct product representation.

I. INTRODUCTION

Covariance estimation is a problem of great interest in many different disciplines, including machine learning, signal processing, economics and bioinformatics. In many applications the number of variables is very large, e.g., in the tens or hundreds of thousands, leading to a number of covariance parameters that

greatly exceeds the number of observations. To address this problem constraints are frequently imposed on the covariance to reduce the number of parameters in the model. For example, the Glasso model of Yuan and Lin [2] and Banerjee *et al* [3] imposes sparsity constraints on the covariance. The Kronecker product model of Dutilleul [4] and Werner *et al* [5] assumes that the covariance can be represented as the Kronecker product of two lower dimensional covariance matrices. The transposable regularized covariance model of Allen *et al* [1] imposes a combination of sparsity and Kronecker product form on the covariance. When there is no missing data, an extension of the alternating optimization algorithm of [4], [5], called the flip flop (FF) algorithm, can be applied to estimate the parameters of this combined sparse and Kronecker product model. In this report we call this algorithm the Kronecker Glasso (KGlasso) and we thoroughly analyze convergence of the algorithm in the high dimensional setting.

As in [5] we assume that there are pf variables whose covariance Σ_0 has the separable positive definite Kronecker product representation:

$$\Sigma_0 = \mathbf{A}_0 \otimes \mathbf{B}_0 \quad (1)$$

where \mathbf{A}_0 is a $p \times p$ positive definite matrix and \mathbf{B}_0 is an $f \times f$ positive definite matrix. This model (1) is relevant to channel modeling for MIMO wireless communications, where \mathbf{A}_0 is a transmit covariance matrix and \mathbf{B}_0 is a receive covariance matrix [6]. The model is also relevant to other transposable models arising in recommendation systems like NetFlix and in gene expression analysis [1].

The Kronecker product Gaussian graphical model has been known for a long time as the matrix normal distribution in the statistics community [7], [4], [8]. Various properties of the matrix variate normal distribution have been studied in [8]. Let us rewrite the problem into matrix form. Consider a $p \times f$ random matrix \mathbf{Z} that follows a matrix normal distribution-i.e. $\mathbf{Z} \sim N_{p,f}(\mathbf{0}; \mathbf{A}_0, \mathbf{B}_0)$ [8]. Then, \mathbf{B}_0 is the row covariance matrix and \mathbf{A}_0 is the column covariance matrix-i.e., $\mathbf{Z}_{:,k} \sim N(\mathbf{0}, [\mathbf{B}_0]_{k,k} \mathbf{A}_0)$ and $\mathbf{Z}_{i,:} \sim N(\mathbf{0}, [\mathbf{A}_0]_{i,i} \mathbf{B}_0)$ ¹. This model further finds applications in geostatistics [9] and genomics [10]. Further applications of matrix-variate normal models include collaborative filtering [11], multi-task learning [12] and face recognition [13]. The Kronecker factorization (1) can easily be generalized to the k -fold case, where $\Sigma_0 = \mathbf{A}_1 \otimes \mathbf{A}_2 \otimes \cdots \otimes \mathbf{A}_k$.

Under the assumption that the measurements are multivariate Gaussian with covariance having the Kronecker product form (1), the maximum likelihood (ML) estimator can be formulated [14]. While the ML estimator has no known closed-form solution, an approximation to the solution can be iteratively

¹Here, $\mathbf{Z}_{i,:}$ is the i th row and $\mathbf{Z}_{:,k}$ is the k th column of the matrix \mathbf{Z} . For concreteness, assume $\mathbf{z} = [\mathbf{z}_1^T, \dots, \mathbf{z}_p^T]^T \sim N(\mathbf{0}, \Sigma_0)$. Then, $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_p]^T$ is the $p \times f$ data matrix with row covariance \mathbf{B}_0 and column covariance \mathbf{A}_0 .

computed via an alternating algorithm: the flip-flop (FF) algorithm [14], [5]. As compared to the standard saturated (unstructured) covariance model, the number of unknown parameters in (1) is reduced from order $\Theta(p^2 f^2)$ to order $\Theta(p^2) + \Theta(f^2)$. This results in a significant reduction in the mean squared error (MSE) and the computational complexity of the maximum likelihood (ML) covariance estimator. This report establishes that further reductions MSE are achievable when the Kronecker matrix factors are known to have sparse inverses, i.e., the measurements obey a sparse Kronecker structured Gaussian graphical model.

The graphical lasso (Glasso) estimator was originally proposed in [2], [3] for estimating a sparse inverse covariance, also called the precision matrix, under an i.i.d. Gaussian observation model. An algorithm for efficiently solving the nonsmooth optimization problem that arises in the Glasso estimator, based on ideas from [3], was proposed in [15]. Glasso has been applied to the time-varying coefficients setting in Zhou *et al* [16] using the kernel estimator for covariances at a target time. Rothman *et al* [17] derived high dimensional convergence rates for a slight variant of Glasso, i.e., only the off-diagonal entries of the estimated precision matrix were penalized using an ℓ_1 -penalty. The high dimensional convergence rate of Glasso was established by Ravikumar *et al* [18]. This report extends their analysis to the case that the covariance has Kronecker structure (1), showing that significantly higher rates of convergence are achievable.

The main contribution is the derivation of the high-dimensional MSE convergence rates for KGlasso as n , p and f go to infinity. When both Kronecker factors are sparse, it is shown that KGlasso *strictly* outperforms FF and Glasso in terms of MSE convergence rate. More specifically, we show KGlasso achieves a convergence rate of $O_P\left(\frac{(p+f)\log\max(p,f,n)}{n}\right)$ and FF achieves a rate of $O_P\left(\frac{(p^2+f^2)\log\max(p,f,n)}{n}\right)$ as $n \rightarrow \infty$, while it is known [17], [16] that Glasso achieves a rate of $O_P\left(\frac{(pf+s)\log\max(p,f,n)}{n}\right)$, where s denotes the number of off-diagonal nonzero elements in the true precision matrix Θ_0 . Simulations show that the performance improvements predicted by the high-dimensional analysis continue to hold for small sample size and moderate matrix dimension. For the example studied in Sec. VIII the empirical MSE of KGlasso is significantly lower than that of Glasso and FF for $p = f = 100$ over the range of n from 10 to 100.

The starting point for the MSE convergence analysis is the large-sample analysis of the FF algorithm (Thm. 1 in [5]). The KGlasso convergence proof uses a large deviation inequality that shows that the dimension of one estimated Kronecker factor, say \mathbf{A} , acts as a multiplier on the number of independent samples when performing inference on the other factor \mathbf{B} . This result is then used to obtain optimal MSE rates in terms of Frobenius norm error between the KGlasso estimated matrix and the ground truth. The

asymptotic MSE convergence analysis is useful since it can be used to guide the selection of sparsity regularization parameters and to determine minimum sample size requirements.

An anonymous reviewer alerted the authors to the related work of Yin and Li [10], published after submission of this paper for publication. Yin and Li obtain high-dimensional MSE bounds for the same matrix normal estimation problem considered here. However, our MSE bounds are tighter than the bounds given in Yin and Li. In particular, neglecting terms of order $\log(pf)$, our bounds are of order $p + f$ as compared to Yin and Li's bounds of order pf , which is significantly weaker for large p, f . We obtain improved bounds due to the use of a tighter concentration inequality, established in Lemma 5.

A. Outline

The outline of the report is as follows. Section II introduces the notation that will be used throughout the report. In Section III, the graphical lasso framework is introduced. Section IV uses this framework to introduce the KGLasso algorithm. Section V shows convergence of KGLasso and characterizes its limit points. The high dimensional MSE convergence rate derivation for the FF algorithm is included in Section VI. Section VII presents a high-dimensional MSE rate result that is used to establish the superiority of KGLasso as compared to FF and standard GLasso, under the sparse Kronecker product representation (1). Section VIII presents simulations that empirically validate the theoretical convergence rates obtained in Section VII.

II. NOTATION

For a square matrix \mathbf{M} , define $|\mathbf{M}|_1 = \|\text{vec}(\mathbf{M})\|_1$ and $|\mathbf{M}|_\infty = \|\text{vec}(\mathbf{M})\|_\infty$, where $\text{vec}(\mathbf{M})$ denotes the vectorized form of \mathbf{M} (concatenation of columns into a vector). $\|\mathbf{M}\|_2$ is the spectral norm of \mathbf{M} . $\mathbf{M}_{i,j}$ and $[\mathbf{M}]_{i,j}$ are the (i,j) th element of \mathbf{M} . Let the inverse transformation (from a vector to a matrix) be defined as: $\text{vec}^{-1}(\mathbf{x}) = \mathbf{X}$, where $\mathbf{x} = \text{vec}(\mathbf{X})$. Define the $pf \times pf$ permutation operator $\mathbf{K}_{p,f}$ such that $\mathbf{K}_{p,f}\text{vec}(\mathbf{N}) = \text{vec}(\mathbf{N}^T)$ for any $p \times f$ matrix \mathbf{N} . For a symmetric matrix \mathbf{M} , $\lambda(\mathbf{M})$ will denote the vector of real eigenvalues of \mathbf{M} and define $\lambda_{\max}(\mathbf{M}) = \|\mathbf{M}\|_2 = \max \lambda_i(\mathbf{M})$ for p.d. symmetric matrix, and $\lambda_{\min}(\mathbf{M}) = \min \lambda_i(\mathbf{M})$. Define the sparsity parameter associated with \mathbf{M} as $s_M = \text{card}(\{(i_1, i_2) : [\mathbf{M}]_{i_1, i_2} \neq 0, i_1 \neq i_2\})$. Let $\kappa(\mathbf{M}) := \frac{\lambda_{\max}(\mathbf{M})}{\lambda_{\min}(\mathbf{M})}$ denote the condition number of a symmetric matrix \mathbf{M} .

For a matrix \mathbf{M} of size $pf \times pf$, let $\{\mathbf{M}(i, j)\}_{i,j=1}^p$ denote its $f \times f$ block submatrices, where each block submatrix is $\mathbf{M}(i, j) = [\mathbf{M}]_{(i-1)f+1:i f, (j-1)f+1:j f}$. Also let $\{\overline{\mathbf{M}}(k, l)\}_{k,l=1}^f$ denote the $p \times p$ block submatrices of the permuted matrix $\overline{\mathbf{M}} = \mathbf{K}_{p,f}^T \mathbf{M} \mathbf{K}_{p,f}$.

Define the set of symmetric matrices $S^p = \{\mathbf{A} \in \mathbb{R}^{p \times p} : \mathbf{A} = \mathbf{A}^T\}$, the set of symmetric positive semidefinite (psd) matrices $S_+^p = \{\mathbf{A} \in \mathbb{R}^{p \times p} : \mathbf{A} = \mathbf{A}^T, \mathbf{z}^T \mathbf{A} \mathbf{z} \geq 0, \forall \mathbf{z} \in \mathbb{R}^p\}$, and the set of symmetric positive definite (pd) matrices $S_{++}^p = \{\mathbf{A} \in \mathbb{R}^{p \times p} : \mathbf{A} = \mathbf{A}^T, \mathbf{z}^T \mathbf{A} \mathbf{z} > 0, \forall \mathbf{z} \neq 0\}$. \mathbf{I}_d is a $d \times d$ identity matrix. It can be shown that S_{++}^p is a convex set, but is not closed [19]. Note that S_{++}^p is simply the interior of the closed convex cone S_+^p .

Statistical convergence rates will be denoted by the $O_P(\cdot)$ notation, which is defined as follows. Consider a sequence of real random variables $\{X_n\}_{n \in \mathbb{N}}$ defined on a probability space (Ω, \mathcal{F}, P) and a deterministic (positive) sequence of reals $\{b_n\}_{n \in \mathbb{N}}$. By $X_n = O_P(1)$ is meant: $\sup_{n \in \mathbb{N}} \Pr(|X_n| > K) \rightarrow 0$ as $K \rightarrow \infty$. The notation $X_n = O_P(b_n)$ is equivalent to $\frac{X_n}{b_n} = O_P(1)$. By $X_n = o_P(1)$ is meant $\Pr(|X_n| > \epsilon) \rightarrow 0$ as $n \rightarrow \infty$ for any $\epsilon > 0$. By $\lambda_n \asymp b_n$ is meant $c_1 \leq \frac{\lambda_n}{b_n} \leq c_2$ for all n , where $c_1, c_2 > 0$ are absolute constants.

III. GRAPHICAL LASSO FRAMEWORK

For simplicity, we assume the number of Kronecker components is $k = 2$. Available are n i.i.d. multivariate Gaussian observations $\{\mathbf{z}_t\}_{t=1}^n$, where $\mathbf{z}_t \in \mathbb{R}^{pf}$, having zero-mean and covariance equal to $\Sigma = \mathbf{A}_0 \otimes \mathbf{B}_0$. Then, the log-likelihood is proportional to:

$$l(\Sigma) := \log \det(\Sigma^{-1}) - \text{tr}(\Sigma^{-1} \hat{\Sigma}_n), \quad (2)$$

where Σ is the positive definite covariance matrix and $\hat{\Sigma}_n = \frac{1}{n} \sum_{t=1}^n \mathbf{z}_t \mathbf{z}_t^T$ is the sample covariance matrix. Recent work [3], [15] has considered ℓ_1 -penalized maximum likelihood estimators for the saturated model where Σ belongs to the unrestricted cone of positive definite matrices. These estimators are known as graphical lasso (Glasso) estimators and are the solution to the ℓ_1 -penalized minimization problem:

$$\hat{\Sigma}_n \in \arg \min_{\Sigma \in S_{++}^p} \{-l(\Sigma) + \lambda |\Sigma^{-1}|_1\}, \quad (3)$$

where $\lambda \geq 0$ is a regularization parameter. If $\lambda > 0$ and $\hat{\Sigma}_n$ is positive definite, then $\hat{\Sigma}_n$ in (3) is the unique minimizer.

A fast iterative algorithm, based on a block coordinate descent approach, exhibiting a computational complexity $\mathcal{O}((pf)^4)$, was developed in [15] to solve the convex program (3). Under the assumption $\lambda \asymp \sqrt{\frac{\log(pf)}{n}}$ solution of (3) was shown to have high dimensional convergence rate [17]:

$$\|\mathbf{G}(\hat{\Sigma}_n, \lambda) - \Theta_0\|_F = O_P\left(\sqrt{\frac{(pf + s) \log(pf)}{n}}\right) \quad (4)$$

where s is an upper bound on the number of non-zero off-diagonal elements of Θ_0 . When $s = O(pf)$, this rate is better than the non-regularized sample covariance estimator:

$$\|\hat{\mathbf{S}}_n - \Sigma_0\|_F = O_P \left(\sqrt{\frac{p^2 f^2}{n}} \right). \quad (5)$$

IV. KRONECKER GRAPHICAL LASSO

Let $\Sigma_0 := \mathbf{A}_0 \otimes \mathbf{B}_0$ denote the true covariance matrix, where $\mathbf{A}_0 := \mathbf{X}_0^{-1}$ and $\mathbf{B}_0 = \mathbf{Y}_0^{-1}$ are the true Kronecker factors. Let \mathbf{A}_{init} denote the initial guess of $\mathbf{A}_0 = \mathbf{X}_0^{-1}$.

Define $J(\mathbf{X}, \mathbf{Y})$ as the negative log-likelihood

$$\begin{aligned} J(\mathbf{X}, \mathbf{Y}) &= \text{tr}((\mathbf{X} \otimes \mathbf{Y})\hat{\mathbf{S}}_n) - f \log \det(\mathbf{X}) \\ &\quad - p \log \det(\mathbf{Y}) \end{aligned} \quad (6)$$

Although the objective (6) is not jointly convex in (\mathbf{X}, \mathbf{Y}) , it is biconvex. This motivates the flip-flop algorithm [4], [5]. Adapting the notation from [5], define the mappings $\hat{\mathbf{A}}(\cdot), \hat{\mathbf{B}}(\cdot)$:

$$\underbrace{\hat{\mathbf{A}}(\mathbf{B})}_{p \times p} = \frac{1}{f} \sum_{k,l=1}^f [\mathbf{B}^{-1}]_{k,l} \overline{\hat{\mathbf{S}}_n}(l, k), \quad (7)$$

$$\underbrace{\hat{\mathbf{B}}(\mathbf{A})}_{f \times f} = \frac{1}{p} \sum_{i,j=1}^p [\mathbf{A}^{-1}]_{i,j} \hat{\mathbf{S}}_n(j, i), \quad (8)$$

where $\overline{\hat{\mathbf{S}}_n} = \mathbf{K}_{p,f}^T \hat{\mathbf{S}}_n \mathbf{K}_{p,f}$ (see Sec. II for definition of $K_{p,f}$). For fixed $\mathbf{B} \in S_{++}^f$, $\hat{\mathbf{A}}(\mathbf{B})$ in (7) is the minimizer of $J(\mathbf{A}^{-1}, \mathbf{B}^{-1})$ over $\mathbf{A} \in S_{++}^p$. A similar interpretation holds for (8). The flip-flop algorithm starts with some arbitrary p.d. matrix \mathbf{A}_{init} and computes \mathbf{B} using (8), then \mathbf{A} using (7), and repeats until convergence. This algorithm does not account for sparsity.

If $\Theta_0 = \mathbf{X}_0 \otimes \mathbf{Y}_0$ is a sparse matrix, which implies that at least one of \mathbf{X}_0 or \mathbf{Y}_0 is sparse, one can penalize the outputs of the flip-flop algorithm and minimize

$$J_\lambda(\mathbf{X}, \mathbf{Y}) = J(\mathbf{X}, \mathbf{Y}) + \bar{\lambda}_X |\mathbf{X}|_1 + \bar{\lambda}_Y |\mathbf{Y}|_1. \quad (9)$$

This leads to an algorithm that we call KGlasso (see Algorithm 1), which sparsifies the Kronecker factors in proportion to the parameters $\bar{\lambda}_X, \bar{\lambda}_Y > 0$.

The Glasso mapping (3) is written as $\mathbf{G}(\cdot, \lambda) : S^d \rightarrow S^d$,

$$\mathbf{G}(\mathbf{T}, \lambda) = \arg \min_{\Theta \in S_{++}^d} \left\{ \text{tr}(\Theta \mathbf{T}) - \log \det(\Theta) + \lambda |\Theta|_1 \right\}. \quad (10)$$

Algorithm 1 Kronecker Graphical Lasso (KGlasso)

- 1: **Input:** $\hat{\mathbf{S}}_n, p, f, n, \bar{\lambda}_X > 0, \bar{\lambda}_Y > 0$
 - 2: **Output:** $\hat{\boldsymbol{\Theta}}_{KGlasso}$
 - 3: Initialize \mathbf{A}_{init} to be positive definite satisfying Assumption 1.
 - 4: $\check{\mathbf{X}} \leftarrow \mathbf{A}_{init}^{-1}$
 - 5: **repeat**
 - 6: $\hat{\mathbf{B}} \leftarrow \frac{1}{p} \sum_{i,j=1}^p [\check{\mathbf{X}}]_{i,j} \hat{\mathbf{S}}_n(j, i)$ (see Eq. (7))
 - 7: $\check{\mathbf{Y}} \leftarrow \mathbf{G}(\hat{\mathbf{B}}, \frac{\bar{\lambda}_Y}{p})$, where $\mathbf{G}(\cdot, \cdot)$ is defined in (10)
 - 8: $\hat{\mathbf{A}} \leftarrow \frac{1}{f} \sum_{k,l=1}^f [\check{\mathbf{Y}}]_{k,l} \overline{\hat{\mathbf{S}}_n}(l, k)$ (see Eq. (8))
 - 9: $\check{\mathbf{X}} \leftarrow \mathbf{G}(\hat{\mathbf{A}}, \frac{\bar{\lambda}_X}{f})$
 - 10: **until** convergence
 - 11: $\hat{\boldsymbol{\Theta}}_{KGlasso} \leftarrow \check{\mathbf{X}} \otimes \check{\mathbf{Y}}$
-

As compared to the $\mathcal{O}(p^4 f^4)$ computational complexity of Glasso, KGlasso has a computational complexity of only $\mathcal{O}(p^4 + f^4)$ ².

V. CONVERGENCE OF KGLASSO ITERATIONS

In this section, we provide an alternative characterization of the KGlasso algorithm and prove convergence to a local minimum of the objective function.

A. Block-Coordinate Reformulation of KGlasso

The KGlasso algorithm can be re-formulated as a block-coordinate optimization of the penalized objective function 9.

- Lemma 1.** 1) Assume $\lambda_X, \lambda_Y \geq 0$ and $\mathbf{X} \in S_{++}^p, \mathbf{Y} \in S_{++}^f$. When one argument of $J_\lambda(\mathbf{X}, \mathbf{Y})$ is fixed, the objective function (9) is convex in the other argument.
- 2) Assume $\hat{\mathbf{S}}_n$ is positive definite. Consider $J_\lambda(\mathbf{X}, \mathbf{Y})$ in (9) with matrix $\mathbf{X} \in S_{++}^p$ fixed. Then, the dual subproblem for minimizing $J_\lambda(\mathbf{X}, \mathbf{Y})$ over \mathbf{Y} is:

$$\max_{|\mathbf{W} - \frac{1}{p} \sum_{i,j=1}^p \mathbf{X}_{i,j} \hat{\mathbf{S}}_n(j, i)|_\infty \leq \lambda_Y} \log \det(\mathbf{W}) \quad (11)$$

²In the sparse Kronecker factor case, this cost can be reduced to $\mathcal{O}(p^3 + f^3)$.

where $\lambda_Y := \bar{\lambda}_Y/p$.

On the other hand, consider (9) with matrix $\mathbf{Y} \in S_{++}^f$ fixed. Then, the dual problem for minimizing $J_\lambda(\mathbf{X}, \mathbf{Y})$ over \mathbf{X} is:

$$\max_{|\mathbf{Z} - \frac{1}{f} \sum_{k,l=1}^f \mathbf{Y}_{k,l} \bar{\hat{\mathbf{S}}}_n(l,k)|_\infty \leq \lambda_X} \log \det(\mathbf{Z}) \quad (12)$$

where $\bar{\hat{\mathbf{S}}}_n := \mathbf{K}_{p,f}^T \hat{\mathbf{S}}_n \mathbf{K}_{p,f}$ and $\lambda_X := \bar{\lambda}_X/f$.

- 3) Strong duality holds for (11) and (12).
- 4) The solutions to (11) and (12) are positive definite.

Proof: See Appendix. ■

Note that both dual subproblems (11) and (12) have a unique solution and the maximum is attained in each one. This follows from the fact that in each case we are maximizing a strictly concave function over a closed convex set. Lemma 1 is similar to the result obtained in [3], but with $(\frac{1}{p} \sum_{i,j=1}^p \mathbf{X}_{i,j} \hat{\mathbf{S}}_n(j,i), \lambda_Y)$ playing the role of $(\hat{\mathbf{S}}_n, \lambda)$, for the “fixed \mathbf{X} ” subproblem.

B. Limit Point Characterization of KGlasso

We will first show that KGlasso converges to a fixed point. Let $J_\lambda(\mathbf{X}, \mathbf{Y})$ be as defined in (9) and define $J_\lambda^{(k)} = J_\lambda(\mathbf{X}^{(k)}, \mathbf{Y}^{(k)})$ for $k = 0, 1, 2, \dots$.

Theorem 1. *If $\hat{\mathbf{S}}_n$ is positive definite, KGlasso converges to a fixed point. Also, we have $J_\lambda^{(k)} \searrow J_\lambda^{(\infty)}$.*

Proof: See Appendix. ■

The following analysis uses Theorem 1 to prove convergence of the KGlasso algorithm to a local minimum. To do this, we consider a more general setting. The KGlasso algorithm is a special case of Algorithm 2. Assuming a k -fold Kronecker product structure for the covariance matrix, the optimization problem (9) can be written in the form:

$$J_\lambda(\mathbf{X}_1, \dots, \mathbf{X}_k) = J_0(\mathbf{X}_1, \dots, \mathbf{X}_k) + \sum_{i=1}^k J_i(\mathbf{X}_i) + \bar{\lambda}_i \eta_1(\mathbf{X}_i) \quad (13)$$

where $\mathbf{X}_i \in S_{++}^{d_i}$, $\eta_1(\mathbf{X}_i) := |\mathbf{X}_i|_1$, $J_0(\mathbf{X}_1, \dots, \mathbf{X}_k) := \text{tr}((\mathbf{X}_1 \otimes \mathbf{X}_2 \otimes \dots \otimes \mathbf{X}_k) \hat{\mathbf{S}}_n)$ and $J_i(\mathbf{X}_i) = -\prod_{i' \neq i} d_{i'} \cdot \log \det(\mathbf{X}_i)$ for $i = 1, \dots, k$.

Without loss of generality, by reshaping matrices into appropriate vectors, (13) can be rewritten as:

$$J_\lambda(\mathbf{x}_1, \dots, \mathbf{x}_k) = J_0(\mathbf{x}_1, \dots, \mathbf{x}_k) + \sum_{i=1}^k J_i(\mathbf{x}_i) + \bar{\lambda}_i \eta_i(\mathbf{x}_i) \quad (14)$$

where the optimization variable is $\mathbf{x} := [\mathbf{x}_1^T, \mathbf{x}_2^T, \dots, \mathbf{x}_k^T]^T \in \mathbb{R}^{d'}$, where $\mathbf{x}_i \in \mathbb{R}^{d_i^2}$ and $d' = \sum_{i=1}^k d_i^2$. For example, $\eta_i(\mathbf{X}_i) = |\mathbf{X}_i|_1 = \|\text{vec}(\mathbf{X}_i)\|_1 = \|\mathbf{x}_i\|_1 = \eta_i(\mathbf{x}_i)$. The mapping $\{J_i\}_{i=0}^k$ can be similarly written in terms of the vectors \mathbf{x}_i instead of the matrices \mathbf{X}_i .

The reader can verify that the objective function (13) satisfies the properties (for $n > pf$) in Appendix D.

The general optimization problem of interest here is:

$$\min_{\mathbf{x} \in \mathbb{R}^{d'}} J_\lambda(\mathbf{x}) \text{ subject to } \text{vec}^{-1}(\mathbf{x}_i) = \mathbf{X}_i \in S_{++}^{d_i}, i = 1, \dots, k \quad (15)$$

The positive definiteness constraints are automatically taken care of by the construction of the algorithm (see Lemma 1.4). Let the dimension of the covariance matrix be denoted by $d := \prod_{i=1}^k d_i$. We assume $n > d$. To solve (15), a block coordinate-descent penalized algorithm is constructed:

Algorithm 2 Block Coordinate-Descent Penalized Algorithm

- 1: **Input:** $\hat{\mathbf{S}}_n, d_i, n, \epsilon > 0, \lambda_i > 0$
 - 2: **Output:** $\hat{\Theta}$
 - 3: Initialize $\mathbf{X}_1^0, \mathbf{X}_2^0, \dots, \mathbf{X}_k^0$ matrices as positive definite matrices, e.g., scaled identity.
 - 4: $\hat{\Theta}_0 \leftarrow \mathbf{X}_1^0 \otimes \mathbf{X}_2^0 \otimes \dots \otimes \mathbf{X}_k^0$
 - 5: $m \leftarrow 0$
 - 6: **repeat**
 - 7: $\hat{\Theta}_{\text{prev}} \leftarrow \hat{\Theta}$
 - 8: $\mathbf{X}_1^m \leftarrow \arg \min_{\mathbf{A}_1 \succ 0} J_\lambda(\mathbf{A}_1, \mathbf{X}_2^{m-1}, \dots, \mathbf{X}_k^{m-1})$
 - 9: $\mathbf{X}_2^m \leftarrow \arg \min_{\mathbf{A}_2 \succ 0} J_\lambda(\mathbf{X}_1^m, \mathbf{A}_2, \dots, \mathbf{X}_k^{m-1})$
 - 10: \vdots
 - 11: $\mathbf{X}_k^m \leftarrow \arg \min_{\mathbf{A}_k \succ 0} J_\lambda(\mathbf{X}_1^m, \mathbf{X}_2^m, \dots, \mathbf{A}_k)$
 - 12: $\hat{\Theta} \leftarrow \mathbf{X}_1^m \otimes \mathbf{X}_2^m \otimes \dots \otimes \mathbf{X}_k^m$
 - 13: $m \leftarrow m + 1$
 - 14: **until** $\|\hat{\Theta}_{\text{prev}} - \hat{\Theta}\| \leq \epsilon$
-

Remark 1. The positive definiteness constraint at each coordinate descent iteration of Algorithms 1 and 2 need not be explicit since the objective function $J_\lambda(\cdot)$ acts as a logarithmic barrier function.

Note that Algorithm 1 is a special case of Algorithm 2. An extension of Theorem 1, assuming $n > d$

or $J_\lambda^* > -\infty$, based on induction, can be used to show that the limit points of the sequence of iterates $(\mathbf{x}^m)_{m \geq 0} = (\mathbf{x}_1^m, \dots, \mathbf{x}_k^m)_{m \geq 0}$ are fixed points.

Remark 2. Note that a necessary condition for \mathbf{x}^* to minimize J_λ is $0 \in \partial J_\lambda(\mathbf{x}^*)$. This is not sufficient however.

We next show that the limit point(s) of $(\mathbf{x}^m)_{m \geq 0}$ are nonempty and are local minima.

Theorem 2. Let $(\mathbf{x}^m) = (\mathbf{x}_1^m, \dots, \mathbf{x}_k^m)_{m \geq 0}$ be a sequence generated by Algorithm 2. Assume $n > d$.

- 1) The algorithm converges to a local minimum.
- 2) If \mathbf{x}^0 is not a local minimum, strict descent follows.

Proof: See Appendix. ■

As a consequence of Theorem 2, we have the following corollary.

Corollary 1. Assuming $n > pf$, the KGlasso algorithm converges to a local minimizer of the objective function (9).

VI. HIGH DIMENSIONAL CONSISTENCY OF FF

In this section, we show that the flip-flop (FF) algorithm achieves the optimal (non-sparse) statistical convergence rate of $O_P\left(\sqrt{\frac{p^2+f^2}{n}}\right)$. This result (see Thm. 3, Thm. 4) allows us to establish that the proposed KGlasso has significantly improved MSE convergence rate (see Thm. 5). We make the following standard assumption on the spectra of the Kronecker factors.

Assumption 1. Uniformly Bounded Spectra

There exist absolute constants $\underline{k}_A, \bar{k}_A, \underline{k}_B, \bar{k}_B, \underline{k}_{A_{init}}, \bar{k}_{A_{init}}$ such that:

- 1a. $0 < \underline{k}_A \leq \lambda_{\min}(\mathbf{A}_0) \leq \lambda_{\max}(\mathbf{A}_0) \leq \bar{k}_A < \infty$
- 1b. $0 < \underline{k}_B \leq \lambda_{\min}(\mathbf{B}_0) \leq \lambda_{\max}(\mathbf{B}_0) \leq \bar{k}_B < \infty$
2. $0 < \underline{k}_{A_{init}} \leq \lambda_{\min}(\mathbf{A}_{init}) \leq \lambda_{\max}(\mathbf{A}_{init}) \leq \bar{k}_{A_{init}} < \infty$

Let $\Sigma_{FF}(3) := \hat{\mathbf{A}}(\hat{\mathbf{B}}(\mathbf{A}_{init})) \otimes \hat{\mathbf{B}}(\hat{\mathbf{A}}(\hat{\mathbf{B}}(\mathbf{A}_{init})))$ denote the 3-step (noniterative) version of the flip-flop algorithm [5]. More generally, let $\Sigma_{FF}(k)$ denote the k -step version of the flip-flop algorithm, and denote its inverse as $\Theta_{FF}(k) = (\Sigma_{FF}(k))^{-1}$.

Theorem 3. Let $\mathbf{A}_0, \mathbf{B}_0, \mathbf{A}_{init}$ satisfy Assumption 1. Let $M = \max(p, f, n)$, $\bar{M} = (\max\{pf^{-1/2} + \sqrt{f}, fp^{-1/2} + \sqrt{p}\})^2$. Assume $\frac{(pf)^3 \bar{M}^2 (\log(M))^3}{(p^2+f^2)^2} = o(n)$ as $p, f, n \rightarrow \infty$. Also, assume $p, f \geq 2$ and $\frac{\bar{M} \sqrt{\log \bar{M}}}{\sqrt{f/p} + \sqrt{p/f}} = O(\sqrt{n})$.

Then, for $k \geq 3$ finite,

$$\|\Theta_{FF}(k) - \Theta_0\|_F = O_P \left(\sqrt{\frac{p^2 + f^2}{n}} \right) \quad (16)$$

as $p, f, n \rightarrow \infty$.

Proof: See Appendix. ■

The next theorem offers a shorter proof of essentially the same convergence rate (16) (up to a logarithm factor) under milder sufficient conditions.

Theorem 4. Let $\mathbf{A}_0, \mathbf{B}_0$, and \mathbf{A}_{init} satisfy Assumption 1 and define $M = \max(p, f, n)$. Assume $p \geq f \geq 2$ and $p \log M \leq C''n$ for some finite constant $C'' > 0$. Finally, assume $n \geq \frac{p}{f} + 1$. Then, for $k \geq 2$ finite,

$$\|\Theta_{FF}(k) - \Theta_0\|_F = O_P \left(\sqrt{\frac{(p^2 + f^2) \log M}{n}} \right) \quad (17)$$

as $n \rightarrow \infty$.

Proof: See Appendix. ■

Remark 3. The sufficient conditions are symmetric with respect to p and f -i.e. for $f \geq p$, the corresponding conditions would become $f \log M \leq C''n$ for some constant $C'' > 0$, and $n \geq \frac{f}{p} + 1$.

For the special case of $p = f$, the sufficient conditions of Thm. 3 become $p^4(\log M)^3 = o(n)$, while the sufficient conditions of Thm. 4 become $p \log M = O(n)$. Thus, to achieve accurate covariance estimation for arbitrarily structured Kronecker factors, the minimal sample size needed is $n = \Omega((p^2 + f^2) \log M)$.

The bound (17) specifies the rate of reduction of the estimation error for the multi-iteration FF algorithm, which includes the three step FF algorithm ($k = 3$) [5] as a special case. The error reduction decreases as long as p and f do not increase too quickly in n .

Note that (17) specifies a faster rate than that of the naive sample covariance matrix estimator (5). Furthermore, since the computational complexity for FF is $\mathcal{O}(p^2 + f^2)$ which is less than the $\mathcal{O}(p^2 f^2)$ complexity of SCM, by exploiting Kronecker structure FF simultaneously achieves improved MSE performance and reduced computational complexity.

VII. HIGH DIMENSIONAL CONSISTENCY OF KGLASSO

In this section, consistency is established for KGLasso as $p, f, n \rightarrow \infty$.




Fig. 1. Root mean square error (RMSE) performance for the flip-flop estimator (FF) (left) and for the standard sample covariance matrix estimator (SCM) (right). SCM performs very poorly in comparison to FF when the covariance matrix decomposes as a Kronecker product. Here, the sample size n is fixed and the dimensions of the Kronecker factors (p, f) vary. Equation (16) is plotted on the left and Equation (5) on the right. Exploiting structure yields a significant reduction in MSE. The magnitude of the colormap reflects the error up to a constant scaling. The colormap in both images is the same, which visually shows the lower RMSE of FF as compared to SCM.

A. MSE convergence rate of KGlasso

Define $\Theta_{KGlasso}(k)$ as the output of the k th compression and sparsification step (two of these steps constitute a full KGlasso iteration).

Theorem 5. Let $\mathbf{A}_0, \mathbf{B}_0, \mathbf{A}_{init}$ satisfy Assumption 1. Let $M = \max(p, f, n)$. Let $\bar{\lambda}_Y^{(1)} \asymp p\sqrt{\frac{\log M}{np}}$ and $\bar{\lambda}_X^{(k)} \asymp \left(\frac{1}{\sqrt{p}} + \frac{1}{\sqrt{f}}\right) f\sqrt{\frac{\log M}{n}}$, $\bar{\lambda}_Y^{(k')} \asymp \left(\frac{1}{\sqrt{p}} + \frac{1}{\sqrt{f}}\right) p\sqrt{\frac{\log M}{n}}$ as $p, f, n \rightarrow \infty$ for all $k \geq 1$ and $k' \geq 2$. Assume sparse \mathbf{X}_0 and \mathbf{Y}_0 , i.e. $s_{X_0} = O(p)$, $s_{Y_0} = O(f)$. Assume $\max\left(\frac{p}{f}, \frac{f}{p}\right) \log M = o(n)$. Then, for $k \geq 2$ finite, we have

$$\|\Theta_{KGlasso}(k) - \Theta_0\|_F = O_P\left(\sqrt{\frac{(p+f)\log M}{n}}\right) \quad (18)$$

as $p, f, n \rightarrow \infty$.

Proof: See Appendix. ■

Theorem 5 offers a strict improvement over standard Glasso [17], [3] and generalizes Thm. 1 in [17] to the case of sparse Kronecker product structure. Thm. 5 generalizes Thm. 4 to the case of sparse Kronecker structure. Comparison between the error expressions (4), (17) and (18) show that, by exploiting both Kronecker structure and sparsity, KGlasso can attain significantly lower estimation error than standard Glasso [17] and FF [5]. To achieve accurate covariance estimation for the sparse Kronecker product model, the minimal sample size needed is $n = \Omega((p+f)\log M)$.

Although Thm. 5 shows a rate on the inverse covariance matrix, this asymptotic rate can be shown to hold for the covariance matrix as well (i.e., the inverse of $\Theta_{KGlasso}$).

Let $\mathbf{B}_1 := \mathbf{G}(\hat{\mathbf{B}}(\mathbf{A}_{init}), \lambda_Y^{(1)})^{-1}$, where \mathbf{G} is defined in (10). Then, $\Theta_{KGlasso}(1) = \mathbf{G}(\hat{\mathbf{A}}(\mathbf{B}_1), \lambda_X^{(1)}) \otimes \mathbf{G}(\hat{\mathbf{B}}(\mathbf{A}_{init}), \lambda_Y^{(1)})$ denotes the KGlasso output after the first two steps of the KGlasso algorithm (or one KGlasso iteration). A graphical depiction of the first three steps of KGlasso is shown in Fig. 2. Define $\mathbf{B}_1 = \mathbf{G}(\hat{\mathbf{B}}(\mathbf{A}_{init}), \lambda_Y^{(1)})^{-1}$, where \mathbf{G} is given in (10). Then, $\Theta_{KGlasso}(1) = \mathbf{G}(\hat{\mathbf{A}}(\mathbf{B}_1), \lambda_X^{(1)}) \otimes \mathbf{G}(\hat{\mathbf{B}}(\mathbf{A}_{init}), \lambda_Y^{(1)})$ denotes the KGlasso output after the first two steps of the KGlasso algorithm (or one KGlasso iteration). Although Thm. 5 shows a rate on the inverse covariance matrix, this asymptotic rate can be shown to hold for the covariance matrix as well (see proof of Thm. 5 in Appendix).




Fig. 2. Illustration of first three iterations of KGlasso. The squares around the blue dots represent the ℓ_∞ balls controlled by the regularization parameter (see dual programs (11) and (12)). As the regularization parameters tend to zero, the balls shrink to the blue points, and KGlasso becomes identical to the FF algorithm.

Figures 3 and 4 graphically compare the MSE convergence rates of KGlasso, FF and standard Glasso as a function of p, f for fixed n . Note that the standard Glasso algorithm would yield an inferior rate to (18) (recall (4)).

The minimal sample size required to achieve accurate covariance estimation is graphically depicted in Fig. 5 for the special case $p = f$. The regions below the lines are the MSE convergence regions-i.e., the MSE convergence rate goes to zero as p, n grow together to infinity at a certain growth rate controlled by these regions. It is shown that KGlasso allows the dimension p to grow almost linearly in n and still achieve accurate covariance estimation (see (18)) and thus, uniformly outperforms FF, Glasso and the




Fig. 3. Root mean square error performance for Kronecker graphical lasso estimator (KGLasso) (left) and flip-flop estimator (FF) (right). FF performs very poorly in comparison to KGLasso when the covariance matrix decomposes as a Kronecker product and both Kronecker factors are sparse. The bound in Equation (18) is plotted on the left and that in Equation (16) on the right. The magnitude of the colormap reflects the error up to a constant scaling.




Fig. 4. Root mean square error performance for Kronecker graphical lasso estimator (KGLasso) (left) and standard Glasso estimator (Glasso) (right). Glasso performs very poorly in comparison to KGLasso when the covariance matrix decomposes as a Kronecker product and both Kronecker factors are sparse. The bound in Equation (18) is plotted on the left and that in Equation (4) on the right. The magnitude of the colormap reflects the error up to a constant scaling.

naive SCM estimators in the case both Kronecker factors are sparse.

B. Discussion

Theorem 5 is established using the large deviation bound in Lemma 5. We provide some intuition on this bound below. Assume that $\mathbf{X}_{init} = \mathbf{X}_0$, or $\mathbf{A}_{init} = \mathbf{X}_{init}^{-1} = \mathbf{A}_0$. Define $\mathbf{W} = \mathbf{X}_0^{1/2} \otimes \mathbf{I}_p$ and $\tilde{\mathbf{z}}_t = \mathbf{W}\mathbf{z}_t$, with i.i.d. $\mathbf{z}_t \sim N(\mathbf{0}, \mathbf{A}_0 \otimes \mathbf{B}_0)$, $t = 1, \dots, n$. Then, $\tilde{\mathbf{z}}_t$ has block-diagonal covariance

$$\text{Cov}(\tilde{\mathbf{z}}_t) = \mathbf{I}_p \otimes \mathbf{B}_0.$$

When \mathbf{W} is applied to the transformed $pf \times pf$ sample covariance matrix, $\hat{\mathbf{S}}_n^W := \mathbf{W}\hat{\mathbf{S}}_n\mathbf{W}^T$, the first step of KGLasso produces an iterate $\hat{\mathbf{Y}}_n^{(1)} = \mathbf{G}(\hat{\mathbf{B}}, \lambda_Y)$ with $\hat{\mathbf{B}} = \frac{1}{p} \sum_{i=1}^p \hat{\mathbf{S}}_n^W(i, i)$ (recall (8)). For suitable $\lambda_Y = \lambda_Y^{(1)}$, $\hat{\mathbf{Y}}_n^{(1)}$ converges to \mathbf{Y}_0 with respect to maximal elementwise norm at a rate $O_P\left(\sqrt{\frac{\log M}{np}}\right)$. The convergence of $\hat{\mathbf{Y}}_n^{(1)}$ is easily established by applying the Chernoff bound and invoking the jointly Gaussian property of the measurements and the block diagonal structure of $\text{Cov}(\tilde{\mathbf{z}}_t)$. Lemma 5 in the




Fig. 5. Graphical depiction of minimal sample size required for KGLasso, FF, GLasso and naive SCM estimators to achieve accurate covariance estimation. The region below the lines constitute the MSE convergence regions-i.e. traveling along a path $(p(n), n)$ as $n \rightarrow \infty$ within such regions implies the MSE convergence rate tends to zero (see (5),(4),(17) and (18)).

Appendix establishes that this rate holds even if $\mathbf{X}_{init} \neq \mathbf{X}_0$ in Assumption 1. In view of the rate of convergence of $\hat{\mathbf{Y}}^{(1)}$, to achieve a reduction in the MSE of \mathbf{Y} , either the sample size n or the dimension p must increase. Lemma 5 provides a tight bound that makes the dependence of the convergence rate explicit in p, f and n . Theorem 5 uses Lemma 5 to show that KGLasso converges to $\mathbf{X}_0 \otimes \mathbf{Y}_0$ with rate $O_P\left(\sqrt{\frac{(p+f) \log M}{n}}\right)$ with respect to Frobenius norm.

VIII. SIMULATION RESULTS

In this section, we empirically validate the convergence rates established in previous sections using Monte Carlo simulation.

Each iteration of the KGLasso involves solving an ℓ_1 penalized covariance estimation problem of dimension 100×100 (Step 6 and Step 8 of KGLasso specified by Algorithm 1). To solve these small sparse covariance estimation problems we used the GLasso algorithm of Hsieh *et al* [20] where the GLasso stopping criterion was determined by monitoring when the duality gap falls below a threshold of 10^{-3} .

To evaluate performance, Monte Carlo simulations were used. Unless otherwise specified, the true matrices $\mathbf{X}_0 := \mathbf{A}_0^{-1}$ and $\mathbf{Y}_0 := \mathbf{B}_0^{-1}$ were unstructured randomly generated positive definite matrices

based on an Erdős-Rényi graph model. First, a square binary matrix \mathbf{C} was generated based on independently and identically distributing “0s” with a probability p^* and “1s” with a probability $1 - p^*$. Then, $\tilde{\mathbf{C}} := (\mathbf{C} + \mathbf{C}^T)/2$ symmetrizes the matrix. The perturbation level ρ was selected as $\rho = 0.05 - \lambda_{\min}(\tilde{\mathbf{C}})$, producing $\mathbf{Y}_0 := \tilde{\mathbf{C}} + \rho \mathbf{I}_f$, the sparse inverse matrix. There was a total of 20 trial runs for each fixed number of samples n . Performance assessment was based on normalized Frobenius norm error in the covariance and precision matrix estimates. The normalized error was calculated using

$$\sqrt{\frac{1}{N_{MC}} \sum_{i=1}^{N_{MC}} \frac{\|\Sigma_0 - \hat{\Sigma}(i)\|_F^2}{\|\Sigma_0\|_F^2}}$$

where N_{MC} is the number of Monte Carlo runs and $\hat{\Sigma}(i)$ is the covariance output from the i th trial run. The same formula can be adapted to calculate the normalized error in the precision matrix $\hat{\Theta}_0$. In the implementation of KGLasso, the regularization parameters were chosen as follows. The initialization was $\mathbf{X}_{init} = \mathbf{I}_p$. The regularization parameters were selected as $\lambda_Y^{(1)} = c_y \sqrt{\frac{\log M}{np}}$, $\lambda_X^{(2)} = c_x \sqrt{\frac{\log M}{nf}} + \lambda_Y^{(1)}$, $\lambda_Y^{(2)} = \lambda_X^{(2)}$, $\lambda_X^{(3)} = \lambda_X^{(2)}$, etc. For Examples 1 and 2 below, the (positive) scaling constants (c_x, c_y) in front of the regularization parameters were chosen experimentally to optimize respective performances. For Example 3, we simply set $c_x = c_y = 0.4$.

A. Example 1

We consider the simple case that \mathbf{X}_0 and \mathbf{Y}_0 are sparse matrices of dimensions $p = 20$ and $f = 10$. Figure 6 shows that $\mathbf{X}_0 \otimes \mathbf{Y}_0$ is a perturbation of \mathbf{I}_{pf} . Figures 7 and 8 compare the root-mean squared error (RMSE) performance in precision and covariance matrices as a function of n . As expected, KGLasso outperforms both naive GLasso and FF over the range of n for both the covariance and the inverse covariance estimation problem. As expected, the FF algorithm suffers in the small sample regime. KGLasso outperforms FF in this regime since it exploits sparsity in addition to Kronecker structure.

B. Example 2

We consider the case when \mathbf{A}_0 is identity and \mathbf{Y}_0 is dense (see Fig. 9). Figures 10 and 11 show similar trends to those exhibited in Figures 7 and 8 for the case that both \mathbf{X}_0 and \mathbf{Y}_0 are sparse.

C. Example 3

We considered the setting where \mathbf{X}_0 and \mathbf{Y}_0 are large sparse matrices of dimension $p = f = 100$ (see Fig. 12). Only 5% of the off-diagonal entries were nonzero for both matrices \mathbf{X}_0 and \mathbf{Y}_0 . The dimension




Fig. 6. Doubly sparse Kronecker matrix representation for simulation example 1. Left panel: left Kronecker factor. Middle panel: right Kronecker factor. Right panel: Kronecker product inverse covariance matrix




Fig. 7. Normalized RMSE of precision matrix estimate $\hat{\Theta} = \hat{\Sigma}^{-1}$ as a function of sample size n for structure exhibited in Fig. 6. KGlasso (Kronecker graphical lasso) uniformly outperforms FF (flip-flop) algorithm and standard Glasso algorithm for all n . Here, $p = 20$ and $f = 10$.




Fig. 8. Normalized RMSE of covariance matrix estimate $\hat{\Sigma}$ as a function of sample size n for structure exhibited in Fig. 6. KGlasso (Kronecker graphical lasso) uniformly outperforms FF (flip-flop) algorithm and standard Glasso algorithm for all n . Here, $p = 20$ and $f = 10$.




Fig. 9. Sparse Kronecker matrix representation for simulation example 2. Left panel: left Kronecker factor. Middle panel: right Kronecker factor. Right panel: Kronecker product inverse covariance matrix.




Fig. 10. Normalized RMSE performance for precision matrix as a function of sample size n . KGLasso (Kronecker graphical lasso) uniformly outperforms FF (flip-flop) algorithm and standard Glasso algorithm for all n . Here, $p = 20$ and $f = 10$.




Fig. 11. Normalized RMSE performance for covariance matrix as a function of sample size n . KGLasso (Kronecker graphical lasso) uniformly outperforms FF (flip-flop) algorithm and standard Glasso algorithm for all n . Here, $p = 20$ and $f = 10$.

of Θ_0 is $d = 10,000$, which was too large for implementation of standard Glasso. Figures 13 and 14 compare the root-mean squared error (RMSE) performance in precision and covariance matrices as a function of n . As expected, KGlasso outperforms both naive Glasso and FF over the range of n for both the covariance and the inverse covariance estimation problem. As expected, the FF algorithm suffers in the small sample regime. KGlasso outperforms FF in this regime since it exploits sparsity in addition to Kronecker structure.

For $n = 10$, there is a 69% (≈ 5.09 dB) RMSE reduction for the precision matrix and 35% RMSE reduction for the covariance matrix when using KGlasso instead of FF. For $n = 100$, there is a 41% (≈ 2.29 dB) RMSE reduction for the precision matrix and 26% RMSE reduction for the covariance matrix. For the small sample regime, there is approximately a 5.09 dB reduction for the precision matrix, which is a significant performance gain.

D. Example 4

Here, the true covariance matrix factors $\mathbf{X}_0 = \mathbf{A}_0^{-1}$ and $\mathbf{Y}_0 = \mathbf{B}_0^{-1}$ were unstructured randomly generated positive definite matrices. First, p random nonzero elements were placed on the diagonal of a square $p \times p$ matrix C . Then, on average p nonzero elements were placed on the off-diagonal and symmetry was imposed. On average, a total of $3p$ elements were nonzero. The resulting matrix $\tilde{\mathbf{C}}$ was regularized to produce the sparse positive definite inverse covariance $\mathbf{Y}_0 = \tilde{\mathbf{C}} + \rho \mathbf{I}_f$, where $\rho = 0.5 - \lambda_{\min}(\tilde{\mathbf{C}})$.

We also compare KGlasso to a natural extension of the FF algorithm that accounts for both sparsity and Kronecker structure. The flip-flop thresholding method (FF/Thres) that we consider consists of first computing the FF solution and then thresholding each estimated precision matrix. To ensure a fair comparison we set the threshold level of FF/Thres that yields exactly the same sparsity factor as the KGlasso estimated precision matrices.

For $n = 10$, there is a 72% (≈ 5.53 dB) RMSE reduction for the precision matrix and 41% RMSE reduction for the covariance matrix when using KGlasso instead of FF. For $n = 10$, there is a 70% (≈ 5.23 dB) RMSE reduction for the precision matrix and 62% RMSE reduction for the covariance matrix when using KGlasso instead of FF/Thres. For $n = 100$, there is a 53% (≈ 3.28 dB) RMSE reduction for the precision matrix and 33% RMSE reduction for the covariance matrix when using KGlasso instead of FF. For $n = 100$, there is a 50% (≈ 3.01 dB) RMSE reduction for the precision matrix and 41% RMSE reduction for the covariance matrix when using KGlasso instead of FF/Thres. For the small sample regime, there is approximately a 5.53 dB reduction for the precision matrix, which is a significant performance gain.




Fig. 12. Sparse Kronecker matrix representation. Left panel: left Kronecker factor. Right panel: right Kronecker factor.




Fig. 13. Normalized RMSE performance for precision matrix as a function of sample size n . KGlasso (Kronecker graphical lasso) uniformly outperforms FF (flip-flop) algorithm for all n . Here, $p = 100$ and $f = 100$. For $n = 10$, there is a 69% RMSE reduction.




Fig. 14. Normalized RMSE performance for covariance matrix as a function of sample size n . KGlasso (Kronecker graphical lasso) uniformly outperforms FF (flip-flop) algorithm for all n . Here, $p = 100$ and $f = 100$. For $n = 10$, there is a 35% RMSE reduction.




Fig. 15. Sparse Kronecker matrix representation. Left panel: left Kronecker factor. Right panel: right Kronecker factor. As the Kronecker-product covariance matrix is of dimension $10,000 \times 10,000$ standard Glasso is not practically implementable for this example.




Fig. 16. Normalized RMSE performance for precision matrix as a function of sample size n . KGLasso (Kronecker graphical lasso) uniformly outperforms FF (flip-flop) algorithm and FF/Thres (flip-flop thresholding) for all n . Here, $p = f = 100$ and $N_{MC} = 40$. The error bars are centered around the mean with \pm one standard deviation. For $n = 10$, there is a 72% RMSE reduction from the FF to KGLasso solution and a 70% RMSE reduction from the FF/Thres to KGLasso.




Fig. 17. Normalized RMSE performance for covariance matrix as a function of sample size n . KGLasso (Kronecker graphical lasso) uniformly outperforms FF (flip-flop) algorithm for all n . Here, $p = f = 100$ and $N_{MC} = 40$. The error bars are centered around the mean with \pm one standard deviation. For $n = 10$, there is a 41% RMSE reduction from the FF to KGLasso solution and a 62% RMSE reduction from the FF/Thres to KGLasso.

We finally remark that the benefit obtained in the reduced convergence rate is not only due to the covariance estimation method chosen, but to the problem it addresses as well-i.e. the assumed true covariance structure.

E. Empirical Rate Comparison

Next, we illustrate the rates obtained in for the dimension setting $p(n) = f(n) = \lceil 8n^\alpha \rceil$, where $\alpha \in \{0.1, 0.2, 0.3\}$. According to the theory developed, for large n , the MSE converges to zero at a certain convergence rate. The predicted rates of FF and KGlasso are fitted on top of the empirical MSE curves by ensuring intersection at $n = 1000$. Fig. 18 shows that the empirical rates match the predicted rates well.




Fig. 18. Precision Matrix MSE convergence as a function of sample size n for FF and KGlasso. The dimensions of the Kronecker factor matrices grow as a function of n as: $p(n) = f(n) = \lceil 8 \cdot n^\alpha \rceil$. The true Kronecker factors were set to identity (so their inverses are fully sparse). The predicted MSE curves according to Thm. 4 and Thm. 5 are also shown. For both KGlasso and FF, the predicted MSE matches the empirical MSE well, thus verifying the rate expressions (17) and (18).

We also show a borderline case $p = f = \lceil n^{0.6} \rceil$. In this case, according to Thm. 4 and Thm. 5, the FF diverges (MSE increases in n), while the KGlasso converges (MSE decreases in n). This is illustrated in Fig. 19. Our predicted rates are plotted on top of the empirical curves.




Fig. 19. Precision Matrix MSE as a function of sample size n for FF and KGLasso. The dimensions of the Kronecker factor matrices grow as a function of n as: $p(n) = f(n) = \lceil n^{0.6} \rceil$. The true Kronecker factors were set to identity (so their inverses are fully sparse). The predicted MSE curves according to Thm. 4 and Thm. 5 are also shown. As predicted by our theory, and by the predicted convergent regions of (n, p) for FF and KGLasso in Fig. 5, the MSE of the FF diverges while the MSE of the KGLasso converges as n increases.

IX. CONCLUSION

We established high dimensional consistency for Kronecker Glasso algorithms that use iterative ℓ_1 -penalized likelihood optimization that exploit both Kronecker structure and sparsity of the covariance. A tight MSE convergence rate was derived for KGLasso, showing significantly better MSE performance than standard Glasso [17], [3] and FF [5]. Simulations validated our theoretical predictions.

As expected, the proposed KGLasso algorithm outperforms other algorithms (Glasso, FF) that do not exploit all prior knowledge about the covariance matrix, i.e., sparsity and Kronecker product structure, that KGLasso exploits. The theory and experiments in this paper establish that this performance gain is substantial, more so as the variable dimension increases. Furthermore, as compared to a simple thresholded FF algorithm, which does account for both sparsity and Kronecker structure, KGLasso has significantly better estimation performance.

ACKNOWLEDGEMENT

The authors thank Prof. Mark Rudelson for very helpful discussions on large deviation theory.

APPENDIX A
PROOF OF LEMMA 1

Proof:

- 1) Let $\theta \in (0, 1)$. Let $\mathbf{X}_1, \mathbf{X}_2 \in S_{++}^p$. Then, by the properties of the Kronecker product and trace:

$$\begin{aligned} & \text{tr}(((\theta\mathbf{X}_1 + (1 - \theta)\mathbf{X}_2) \otimes \mathbf{Y})\hat{\mathbf{S}}_n) \\ &= \theta \text{tr}((\mathbf{X}_1 \otimes \mathbf{Y})\hat{\mathbf{S}}_n) + (1 - \theta) \text{tr}((\mathbf{X}_2 \otimes \mathbf{Y})\hat{\mathbf{S}}_n) \end{aligned}$$

The function $g(\mathbf{X}_1) := -\log \det(\mathbf{X}_1)$ is a convex function in \mathbf{X}_1 over the set S_{++}^p [19]. By the triangle inequality:

$$|\theta\mathbf{X}_1 + (1 - \theta)\mathbf{X}_2|_1 \leq \theta|\mathbf{X}_1|_1 + (1 - \theta)|\mathbf{X}_2|_1$$

Finally, the sum of convex functions is convex. The set S_{++}^p is a convex set for any $p \in \mathbb{N}$. The other half of the argument follows by symmetry.

- 2) By symmetry we only need prove that (12) is the dual of $\min_{\mathbf{Y} \in S_{++}^f} J_\lambda(\mathbf{X}, \mathbf{Y})$. By standard duality relations between ℓ_1 and ℓ_∞ norms [19] and symmetry of \mathbf{Y} :

$$|\mathbf{Y}|_1 = \max_{\mathbf{U} \in S^f: |\mathbf{U}|_\infty \leq 1} \text{tr}(\mathbf{Y}\mathbf{U})$$

The maximum is attained at $\mathbf{U}_{i,j} = \frac{\mathbf{Y}_{i,j}}{|\mathbf{Y}_{i,j}|}$ for $\mathbf{Y}_{i,j} \neq 0$ and at $\mathbf{U}_{i,j} = 0$ for $\mathbf{Y}_{i,j} = 0$. Using this in (9) and invoking the saddlepoint inequality:

$$\begin{aligned} & \min_{\mathbf{Y} \in S_{++}^f} \text{tr}((\mathbf{X} \otimes \mathbf{Y})\hat{\mathbf{S}}_n) - p \log \det(\mathbf{Y}) + p\lambda_Y |\mathbf{Y}|_1 \\ &= \min_{\mathbf{Y} \in S_{++}^f} \max_{|\mathbf{U}|_\infty \leq \lambda_Y} \left\{ \text{tr}((\mathbf{X} \otimes \mathbf{Y})\hat{\mathbf{S}}_n) - p \log \det(\mathbf{Y}) \right. \\ & \quad \left. + p \text{tr}(\mathbf{Y}\mathbf{U}) \right\} \\ &\geq \max_{|\mathbf{U}|_\infty \leq \lambda_Y} \min_{\mathbf{Y} \in S_{++}^f} \left\{ \text{tr}((\mathbf{X} \otimes \mathbf{Y})\hat{\mathbf{S}}_n) - p \log \det(\mathbf{Y}) \right. \\ & \quad \left. + p \text{tr}(\mathbf{Y}\mathbf{U}) \right\} \end{aligned} \tag{19}$$

When the equality in (19) is achieved, (\mathbf{U}, \mathbf{Y}) is a saddlepoint and the duality gap is zero. Rewrite the objective function, denoted $\tilde{J}_\lambda(\cdot, \cdot)$, in the minimax operation (19):

$$\tilde{J}_\lambda(\mathbf{X}, \mathbf{Y}) := \text{tr}((\mathbf{X} \otimes \mathbf{Y})(\hat{\mathbf{S}}_n + \tilde{\mathbf{U}}(\mathbf{X}))) - p \log \det(\mathbf{Y})$$

where $\tilde{\mathbf{U}}(\mathbf{X}) = p \frac{\mathbf{I}_p \otimes \mathbf{U}}{\text{tr}(\mathbf{X})}$. Define $\mathbf{M} = \hat{\mathbf{S}}_n + \tilde{\mathbf{U}}(\mathbf{X})$. To evaluate $\min_{\mathbf{Y} \in S_{++}^f} \tilde{J}_\lambda(\mathbf{X}, \mathbf{Y})$ in (19), we invoke the KKT conditions to obtain the solution $\mathbf{Y} = \left(\frac{1}{p} \sum_{i,j=1}^p \mathbf{X}_{i,j} \mathbf{M}(j, i) \right)^{-1}$. Define

$\mathbf{W} = \mathbf{Y}^{-1}$ as the dual space variable. Using this in (19):

$$\max_{|\mathbf{W} - \frac{1}{p} \sum_{i,j=1}^p \mathbf{X}_{i,j} \hat{\mathbf{S}}_n(j,i)|_\infty \leq \lambda_Y} \{p \log \det(\mathbf{W}) + pf\} \quad (20)$$

where the constraint set was obtained in terms of \mathbf{W} by observing that $\tilde{\mathbf{U}}(\mathbf{X})(j,i) = \frac{p\mathbf{U}}{\text{tr}(\mathbf{X})} I(j=i)$, and $I(\cdot)$ is the indicator function. It is evident that (20) is equivalent to (11).

- 3) It suffices to verify that the duality induced by the saddle point formulation is equivalent to Lagrangian duality (see Section 5.4 in [19]). Slater's constraint qualification (see Section 5.3.2 in [19]) trivially holds for the convex problem $\min_{\mathbf{Y} \in S_{++}^f} J_\lambda(\mathbf{X}, \mathbf{Y})$ and the corresponding convex problem $\min_{\mathbf{Y} \in S_{++}^f} \tilde{J}_\lambda(\mathbf{X}, \mathbf{Y})$. Since the objective function of each dual problem has an optimal objective that is bounded below, Slater's constraint qualification also implies that the dual optimal solution is attained.
- 4) From [5], it follows that if $\hat{\mathbf{S}}_n$ is p.d., each "compression step" (see lines 6 and 8 in Algorithm 1) yields a p.d. matrix. Combining this with the positive definiteness of the Glasso estimator [3], we conclude that the first subiteration of KGlasso yields a p.d. matrix. A simple induction, combined with the fact that the Kronecker product of p.d. matrices is p.d., establishes that (11) and (12) are p.d.

■

APPENDIX B

PROOF OF THEOREM 1

Proof: Recall that the basic optimization problem (3) is

$$\min_{\mathbf{X} \in S_{++}^p, \mathbf{Y} \in S_{++}^f} J_\lambda(\mathbf{X}, \mathbf{Y})$$

Let $J^* := \inf_{\mathbf{X} \in S_{++}^p, \mathbf{Y} \in S_{++}^f} J_\lambda(\mathbf{X}, \mathbf{Y})$ be the optimal primal value. Note that $J_\lambda^* > -\infty$ when $\hat{\mathbf{S}}_n \in S_{++}^{pf}$. Now, consider the first step in Algorithm 1. Fix $\mathbf{X} = \mathbf{X}^{(k-1)}$ and optimize over $\mathbf{Y} \in S_{++}^f$. Invoking Lemma 1, we have $\mathbf{Y}^{(k)} = \arg \min_{\mathbf{Y} \in S_{++}^f} J_\lambda(\mathbf{X}^{(k-1)}, \mathbf{Y})$. Note, by induction $\mathbf{Y}^{(k)}$ remains positive definite, assuming $\hat{\mathbf{S}}_n$ is positive definite and $\mathbf{X}^{(0)}$ is positive definite. Considering the second step in Algorithm 1, we fix $\mathbf{Y} = \mathbf{Y}^{(k)}$ and obtain $\mathbf{X}^{(k)} = \arg \min_{\mathbf{X} \in S_{++}^p} J_\lambda(\mathbf{X}, \mathbf{Y}^{(k)})$, so that

$$J_\lambda(\mathbf{X}^{(k)}, \mathbf{Y}^{(k)}) \leq J_\lambda(\mathbf{X}^{(k-1)}, \mathbf{Y}^{(k)}) \leq J_\lambda(\mathbf{X}^{(k-1)}, \mathbf{Y}^{(k-1)}) \quad (21)$$

By induction on the number of iterations of the penalized flip-flop algorithm, we conclude that the iterates yield a nonincreasing sequence of objective functions. Since $\lambda_X |\mathbf{X}|_1, \lambda_Y |\mathbf{Y}|_1 \geq 0$, we see that

the maximum likelihood objective function provides a lower bound to the optimal primal value

$$J_\lambda(\mathbf{X}_{K\text{Glasso}}, \mathbf{Y}_{K\text{Glasso}}) \geq J_\lambda^* \geq J_\lambda(\mathbf{X}_{ML}, \mathbf{Y}_{ML}) > -\infty \quad (22)$$

Thus, the sequence $\{J_\lambda^{(k)} : k \geq 0\}$ forms a nonincreasing sequence bounded below (since for $n > pf$, the log-likelihood function is bounded above by the log-likelihood evaluated at the sample mean and sample covariance matrix). The monotone convergence theorem for sequences [21] implies that $\{J_\lambda^{(k)}\}$ converges monotonically to $J_\lambda^{(\infty)} = \inf_k J_\lambda^{(k)}$. By the alternating minimization, we conclude that the sequence of iterates $\{(\mathbf{X}^{(k)}, \mathbf{Y}^{(k)})\}_k$ converges since the minimizer at each Glasso step is unique. ■

APPENDIX C

SUBDIFFERENTIAL CALCULUS REVIEW

As sparse Kronecker Glasso involves non-smooth objective functions, we review a few definitions and facts from subdifferential calculus [22].

Definition 1. By J -attentive convergence denoted as, $\mathbf{x}^n \xrightarrow{J} \mathbf{x}$, we mean that: $\mathbf{x}^n \rightarrow \mathbf{x}$ with $J(\mathbf{x}^n) \rightarrow J(\mathbf{x})$ as $n \rightarrow \infty$.

The role of J -attentive convergence is to make sure that subgradients at a point $\bar{\mathbf{x}}$ reflect no more than the local geometry of $\text{epi}(J)$ around $(\bar{\mathbf{x}}, J(\bar{\mathbf{x}}))$.

Definition 2. Consider a proper lower semicontinuous (LSC) function $g : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$. Let $\bar{\mathbf{x}}$ be such that $J(\bar{\mathbf{x}}) < \infty$.

For $\mathbf{v} \in \mathbb{R}^d$,

- a) \mathbf{v} is a regular subgradient of J at $\bar{\mathbf{x}}$ (i.e., $\mathbf{v} \in \hat{\partial}J(\bar{\mathbf{x}})$) if $\liminf_{\mathbf{x} \neq \bar{\mathbf{x}}, \mathbf{x} \rightarrow \bar{\mathbf{x}}} \frac{J(\mathbf{x}) - J(\bar{\mathbf{x}}) - \mathbf{v}^T(\mathbf{x} - \bar{\mathbf{x}})}{\|\mathbf{x} - \bar{\mathbf{x}}\|} \geq 0$.
- b) \mathbf{v} is a general subgradient of J at $\bar{\mathbf{x}}$ (i.e., $\mathbf{v} \in \partial J(\bar{\mathbf{x}})$) if there exists subsequences $\mathbf{x}^n \xrightarrow{J} \bar{\mathbf{x}}$ and $\mathbf{v}^n \in \hat{\partial}J(\mathbf{x}^n)$ such that $\mathbf{v}^n \rightarrow \mathbf{v}$.

Let $\bar{\mathbf{x}}$ be such that $J(\bar{\mathbf{x}}) < \infty$. It can be shown that $\partial J(\bar{\mathbf{x}}) = \limsup_{\mathbf{x} \rightarrow \bar{\mathbf{x}}} \hat{\partial}J(\mathbf{x})$, $\hat{\partial}J(\bar{\mathbf{x}}) \subset \partial J(\bar{\mathbf{x}})$ and both sets are closed.

Define the set of critical points $C_J := \{\mathbf{x} : 0 \in \partial J(\mathbf{x})\} = C_{J,\min} \cup C_{J,\text{saddle}} \cup C_{J,\max}$, where $C_{J,\min}$ contains all the local minima, $C_{J,\text{saddle}}$ contains all the saddle points and $C_{J,\max}$ contains all the local maxima.

Definition 3. Let $A \subseteq \mathbb{R}^n$. Define the distance from a point $\mathbf{x} \in \mathbb{R}^n$ to the set A as $d(\mathbf{x}, A) := \inf_{\mathbf{a} \in A} \|\mathbf{x} - \mathbf{a}\|_2$.

APPENDIX D

PROPERTIES OF OBJECTIVE FUNCTION J_λ

The following set of properties will be used in Lemmas 2, 3 and Theorem 2.

Property 1. 1. $J_0 : \mathbb{R}^d \rightarrow \mathbb{R}$ is continuously differentiable (i.e., $f_0 \in C^1$)

2. $\nabla J_0 : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is uniformly continuous on bounded subsets $B \subset \mathbb{R}^d$

3. $J_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R} \cup \{+\infty\}$ is proper³ and lower semicontinuous (LSC), for $i = 1, \dots, k$

4. $\eta_i : \mathbb{R}^d \rightarrow \mathbb{R}_+$ is uniformly continuous and bounded on bounded subsets $B \subset \mathbb{R}^d$, for $i = 1, \dots, k$

5. J_λ is bounded below-i.e. $J_\lambda^* > -\infty$

6. J_λ is strictly convex in at least one block (for all the rest of the blocks held fixed)

where $J_\lambda^* = \inf_{\mathbf{x}_i \in S_{++}^{d_i}} J_\lambda(\mathbf{X}_1, \dots, \mathbf{X}_k)$ is the optimal primal value.

APPENDIX E

LEMMA 2

Lemma 2. Given the notation established in Definition 2 and J_λ given by (14), we have:

$$\begin{aligned} \partial J_\lambda(\mathbf{x}_1, \dots, \mathbf{x}_k) &= \times_{i=1}^k \{ \nabla_{\mathbf{x}_i} J_0(\mathbf{x}_1, \dots, \mathbf{x}_k) + \partial J_i(\mathbf{x}_i) \\ &\quad + \bar{\lambda}_i \partial \eta_i(\mathbf{x}_i) \} \\ &= \times_{i=1}^k \{ \partial_{\mathbf{x}_i} J_\lambda(\mathbf{x}_1, \dots, \mathbf{x}_k) \} \end{aligned} \quad (23)$$

where $\partial_{\mathbf{x}_i} J_\lambda(\mathbf{x}_1, \dots, \mathbf{x}_k)$ is the partial differential operator while all $\{\mathbf{x}_j : j \neq i\}$ are held fixed.

Proof: First note that we have:

$$\begin{aligned} \partial J_\lambda(\mathbf{x}_1, \dots, \mathbf{x}_k) &= \nabla J_0(\mathbf{x}_1, \dots, \mathbf{x}_k) + \partial \left\{ \sum_{i=1}^k J_i(\mathbf{x}_i) \right. \\ &\quad \left. + \sum_{i=1}^k \bar{\lambda}_i \eta_i(\mathbf{x}_i) \right\} \end{aligned} \quad (24)$$

$$= \nabla J_0(\mathbf{x}_1, \dots, \mathbf{x}_k) + \partial \left\{ \sum_{i=1}^k J_i(\mathbf{x}_i) \right\} + \partial \left\{ \sum_{i=1}^k \bar{\lambda}_i \eta_i(\mathbf{x}_i) \right\} \quad (25)$$

$$= \nabla J_0(\mathbf{x}_1, \dots, \mathbf{x}_k) + \times_{i=1}^k \{ \partial J_i(\mathbf{x}_i) \} + \times_{i=1}^k \{ \bar{\lambda}_i \partial \eta_i(\mathbf{x}_i) \} \quad (26)$$

$$= \times_{i=1}^k \{ \nabla_{\mathbf{x}_i} J_0(\mathbf{x}_1, \dots, \mathbf{x}_k) + \partial J_i(\mathbf{x}_i) + \bar{\lambda}_i \partial \eta_i(\mathbf{x}_i) \} \quad (27)$$

³A function $J : \mathbb{X} \rightarrow \mathbb{R} \cup \{\pm\infty\}$ is proper if $\text{dom}(J) = \{x \in \mathbb{X} : J(x) < \infty\} \neq \emptyset$ and $J(x) > -\infty, \forall x \in \mathbb{X}$.

where (24) follows from Property 1 and Exercise 8.8(c) in [22], (25) follows from Corollary 10.9 in [22], (26) follows from Proposition 10.5 and Equation 10(6) p.438 in [22] since $\lambda_i > 0$, and finally (27) follows from Minkowski sum properties. ■

APPENDIX F

LEMMA 3

Lemma 3. *Let m denote the iteration index. For $m \in \mathbb{N}$, define:*

$$\begin{aligned}
 (\mathbf{x}_1^m)^\circ &:= \nabla_{\mathbf{x}_1} J_0(\mathbf{x}_1^m, \mathbf{x}_2^m, \dots, \mathbf{x}_k^m) \\
 &\quad - \nabla_{\mathbf{x}_1} J_0(\mathbf{x}_1^m, \mathbf{x}_2^{m-1}, \dots, \mathbf{x}_k^{m-1}) \\
 (\mathbf{x}_2^m)^\circ &:= \nabla_{\mathbf{x}_2} J_0(\mathbf{x}_1^m, \mathbf{x}_2^m, \dots, \mathbf{x}_k^m) \\
 &\quad - \nabla_{\mathbf{x}_2} J_0(\mathbf{x}_1^m, \mathbf{x}_2^m, \mathbf{x}_3^{m-1}, \dots, \mathbf{x}_k^{m-1}) \\
 &\quad \vdots \\
 (\mathbf{x}_j^m)^\circ &:= \nabla_{\mathbf{x}_j} J_0(\mathbf{x}_1^m, \mathbf{x}_2^m, \dots, \mathbf{x}_k^m) \\
 &\quad - \nabla_{\mathbf{x}_j} J_0(\mathbf{x}_1^m, \dots, \mathbf{x}_j^m, \mathbf{x}_{j+1}^{m-1}, \dots, \mathbf{x}_k^{m-1}) \\
 &\quad \vdots \\
 (\mathbf{x}_k^m)^\circ &:= 0
 \end{aligned}$$

Then, $((\mathbf{x}_1^m)^\circ, \dots, (\mathbf{x}_k^m)^\circ) \in \partial J_\lambda(\mathbf{x}_1^m, \dots, \mathbf{x}_k^m)$. Also, for all convergent subsequences $(\mathbf{x}^{m_j})_j$ of the sequence $(\mathbf{x}^m)_m$, we have

$$d(0, \partial J_\lambda(\mathbf{x}_1^{m_j}, \dots, \mathbf{x}_k^{m_j})) \rightarrow 0 \text{ as } j \rightarrow \infty$$

Proof: From Algorithm 2, we have:

$$\begin{aligned}
 \mathbf{x}_1^m &\in \arg \min_{\mathbf{x}_1} J_\lambda(\mathbf{x}_1, \mathbf{x}_2^{m-1}, \dots, \mathbf{x}_k^{m-1}) \\
 \mathbf{x}_2^m &\in \arg \min_{\mathbf{x}_2} J_\lambda(\mathbf{x}_1^m, \mathbf{x}_2, \mathbf{x}_3^{m-1}, \dots, \mathbf{x}_k^{m-1}) \\
 &\quad \vdots \\
 \mathbf{x}_k^m &\in \arg \min_{\mathbf{x}_k} J_\lambda(\mathbf{x}_1^m, \dots, \mathbf{x}_{k-1}^m, \mathbf{x}_k)
 \end{aligned}$$

The first subiteration step of the algorithm implies that $0 \in \partial_{\mathbf{x}_1} J_\lambda(\mathbf{x}_1^m, \mathbf{x}_2^{m-1}, \dots, \mathbf{x}_k^{m-1})$, the second subiteration step implies $0 \in \partial_{\mathbf{x}_2} J_\lambda(\mathbf{x}_1^m, \mathbf{x}_2^m, \mathbf{x}_3^{m-1}, \dots, \mathbf{x}_k^{m-1})$, etc. Rewriting these using Lemma 2, we

have:

$$\begin{aligned}
0 &\in \nabla_{\mathbf{x}_1} J_0(\mathbf{x}_1^m, \mathbf{x}_2^{m-1}, \dots, \mathbf{x}_k^{m-1}) + \partial J_1(\mathbf{x}_1^m) + \bar{\lambda}_1 \partial \eta_1(\mathbf{x}_1^m) \\
0 &\in \nabla_{\mathbf{x}_2} J_0(\mathbf{x}_1^m, \mathbf{x}_2^m, \mathbf{x}_3^{m-1}, \dots, \mathbf{x}_k^{m-1}) + \partial J_2(\mathbf{x}_2^m) + \bar{\lambda}_2 \partial \eta_2(\mathbf{x}_2^m) \\
&\vdots \\
0 &\in \nabla_{\mathbf{x}_k} J_0(\mathbf{x}_1^m, \mathbf{x}_2^m, \dots, \mathbf{x}_k^m) + \partial J_k(\mathbf{x}_k^m) + \bar{\lambda}_k \partial \eta_k(\mathbf{x}_k^m)
\end{aligned}$$

This implies that for $i = 1, \dots, k$:

$$(\mathbf{x}_i^m)^\circ \in \nabla_{\mathbf{x}_i} J_0(\mathbf{x}_1^m, \mathbf{x}_2^m, \dots, \mathbf{x}_k^m) + \partial J_i(\mathbf{x}_i^m) + \bar{\lambda}_i \partial \eta_i(\mathbf{x}_i^m)$$

It is important to note that $\partial \eta_i(\mathbf{x}) \neq \emptyset, \forall \mathbf{x} \in \mathbb{R}^{d_i}$, for $i = 1, \dots, k$, as a result of property 1.4. To see why, apply Corollary 8.10 in [22] since η_i is finite and locally LSC at every point in its domain. This in turn implies $((\mathbf{x}_1^m)^\circ, \dots, (\mathbf{x}_k^m)^\circ) \in \partial J_\lambda(\mathbf{x}_1^m, \dots, \mathbf{x}_k^m)$ by Lemma 3.

Now, take an arbitrary convergent subsequence $(\mathbf{x}_1^{m_j}, \dots, \mathbf{x}_k^{m_j})_j$ of $(\mathbf{x}_1^m, \dots, \mathbf{x}_k^m)_m$. The convergence of $(\mathbf{x}_1^{m_j}, \dots, \mathbf{x}_k^{m_j})_j$ implies the convergence of $(\mathbf{x}_1^{m_j}, \mathbf{x}_2^{m_j-1}, \dots, \mathbf{x}_k^{m_j-1})_j$, and $(\mathbf{x}_1^{m_j}, \dots, \mathbf{x}_i^{m_j}, \mathbf{x}_{i+1}^{m_j-1}, \dots, \mathbf{x}_k^{m_j-1})_j$ for $i = 2, \dots, k-1$. Taking $j \rightarrow \infty$ and using properties 1.2, we see that $\lim_{j \rightarrow \infty} d(0, \partial J_\lambda(\mathbf{x}_1^{m_j}, \dots, \mathbf{x}_k^{m_j})) = 0$ since $\lim_{j \rightarrow \infty} ((\mathbf{x}_1^{m_j})^\circ, \dots, (\mathbf{x}_k^{m_j})^\circ) = (0, \dots, 0)$. ■

APPENDIX G

PROOF OF THEOREM 2

Proof:

- 1) Let $L(\mathbf{x}^0) = L(\mathbf{x}_1^0, \dots, \mathbf{x}_k^0)$ be the set of all limit points of $(\mathbf{x}^m)_{m \geq 0}$ starting from \mathbf{x}^0 . The block-coordinate descent algorithm, Algorithm 2, implies

$$\begin{aligned}
&J_0(\mathbf{x}_1^m, \mathbf{x}_2^{m-1}, \dots, \mathbf{x}_k^{m-1}) + J_1(\mathbf{x}_1^m) + \bar{\lambda}_1 \eta_1(\mathbf{x}_1^m) \\
&\leq J_0(\alpha_1, \mathbf{x}_2^{m-1}, \dots, \mathbf{x}_k^{m-1}) + J_1(\alpha_1) + \bar{\lambda}_1 \eta_1(\alpha_1)
\end{aligned}$$

for any $\alpha_1 \in \mathbb{R}^{d_1}$. Now, assume there exists a subsequence $(\mathbf{x}^{m_j})_j$ of $(\mathbf{x}^m)_m$ that converges to \mathbf{x}^* , where \mathbf{x}^* is a limit point. This implies that $(\mathbf{x}_1^{m_j}, \mathbf{x}_2^{m_j-1}, \dots, \mathbf{x}_k^{m_j-1}) \rightarrow \mathbf{x}^*$ as $j \rightarrow \infty$. The above inequality combined with properties 1.1 and 1.4 (i.e. the continuity J_0 and η_i) then implies that

$$\begin{aligned}
\limsup_{j \rightarrow \infty} J_1(\mathbf{x}_1^{m_j}) + J_0(\mathbf{x}_1^*, \dots, \mathbf{x}_k^*) &\leq J_1(\alpha_1) \\
&+ J_0(\alpha_1, \mathbf{x}_2^*, \dots, \mathbf{x}_k^*) + \bar{\lambda}_1 (\eta_1(\alpha_1) - \eta_1(\mathbf{x}_1^*))
\end{aligned}$$

for all $\alpha_1 \in \mathbb{R}^{d_1^2}$. Taking $\alpha_1 = \mathbf{x}_1^*$ then yields $\limsup_{j \rightarrow \infty} J_1(\mathbf{x}_1^{m_j}) \leq J_1(\mathbf{x}_1^*)$. Using the lower semi-continuity property of J_1 (property 1.3), we have $\liminf_{j \rightarrow \infty} J_1(\mathbf{x}_1^{m_j}) \geq J_1(\mathbf{x}_1^*)$. Thus, $\lim_{j \rightarrow \infty} J_1(\mathbf{x}_1^{m_j}) = J_1(\mathbf{x}_1^*)$. By a similar line of reasoning, it can be shown that $J_i(\mathbf{x}_i^{m_j}) \rightarrow J_i(\mathbf{x}_i^*)$ as $j \rightarrow \infty$, for $i = 1, \dots, k$. As a result, $\sum_{i=1}^k J_i(\mathbf{x}_i^{m_j}) \rightarrow \sum_{i=1}^k J_i(\mathbf{x}_i^*)$ as $j \rightarrow \infty$. Since $J_0(\cdot)$ is jointly continuous, $J_0(\mathbf{x}_1^{m_j}, \dots, \mathbf{x}_k^{m_j}) \rightarrow J_0(\mathbf{x}_1^*, \dots, \mathbf{x}_k^*)$. By continuity of $\eta_i(\cdot)$, $\sum_{i=1}^k \bar{\lambda}_i \eta_i(\mathbf{x}_i^{m_j}) \rightarrow \sum_{i=1}^k \bar{\lambda}_i \eta_i(\mathbf{x}_i^*)$. Thus, $J_\lambda(\mathbf{x}^{m_j}) \rightarrow J_\lambda(\mathbf{x}^*)$ as $j \rightarrow \infty$.

Now, Lemma 3 implies that $((\mathbf{x}^{m_j})^\circ) \in \partial J_\lambda(\mathbf{x}^{m_j})$. Since the subsequence $(\mathbf{x}^{m_j})_j$ is convergent, by Lemma 3, we have $(\mathbf{x}^{m_j})^\circ \rightarrow 0$ as $j \rightarrow \infty$. As a result, since $\partial J_\lambda(\mathbf{x}^{m_j})$ is closed (see Theorem 8.6 in [22]) for all j , we conclude that $\mathbf{x}^* \in C_J$. Thus, $L(\mathbf{x}^0) \subseteq C_J$.

We have thus proved that limit points are critical points of the objective function.

We can rule out convergence to local maxima thanks to property 1.6. Let us show this rigorously. Assume there exists a local maximum at $\mathbf{x}' = (\mathbf{x}'_1, \dots, \mathbf{x}'_k)$. Then, there exists $r > 0$ such that $J_\lambda(\mathbf{x}) \leq J_\lambda(\mathbf{x}')$ for all \mathbf{x} such that $\|\mathbf{x} - \mathbf{x}'\|_2 < r$. Fix $\mathbf{x}_i = \mathbf{x}'_i$ for all $i \neq 1$. Without loss of generality, assume J_λ is strictly convex in the first block. Since strict convexity is maintained through linear transformation, without loss of generality, assume $d_1 = 1$. Let $\epsilon < r$. Define $x_{1,\epsilon} = x'_1 - \epsilon$ and $x_{2,\epsilon} = x'_1 + \epsilon$. Define $x_\theta = \theta x_{1,\epsilon} + (1 - \theta)x_{2,\epsilon}$, where $\theta \in (0, 1)$. Since $\|[x_\theta; \mathbf{x}_{\neq 1}] - \mathbf{x}'\|_2 = |x_\theta - x'_1| = \epsilon(1 - 2\theta) < r$, by the local maximum definition, there exists $\epsilon \in (0, r)$ small enough such that

$$\theta J_\lambda(x_{1,\epsilon}, \mathbf{x}'_{\neq 1}) + (1 - \theta) J_\lambda(x_{2,\epsilon}, \mathbf{x}'_{\neq 1}) \leq J_\lambda(x_\theta, \mathbf{x}'_{\neq 1})$$

for some $\theta \in (0, 1)$. Since $\epsilon > 0$, we have $x_{1,\epsilon} \neq x_{2,\epsilon}$, and this contradicts strict convexity. Thus, there are no local maxima. ⁴

Next, we use the non-existence of local maxima and continuity of J_λ to rule out convergence to saddle points. Assume there exists a saddlepoint at \mathbf{x}_s . Then, by definition, $0 \in J_\lambda(\mathbf{x}_s)$ and \mathbf{x}_s is not a local maximum or a local minimum. Since \mathbf{x}_s is not a local minimum, for all $\epsilon > 0$, there exists a point \mathbf{x}' such that $\|\mathbf{x}' - \mathbf{x}_s\|_2 < \epsilon$ and $J_\lambda(\mathbf{x}_s) > J_\lambda(\mathbf{x}')$. By continuity, it follows that there exists $\delta > 0$ such that for all \mathbf{x} satisfying $\|\mathbf{x} - \mathbf{x}'\|_2 < \delta$, we have $J_\lambda(\mathbf{x}_s) > J_\lambda(\mathbf{x})$, which implies that \mathbf{x}_s is a local maximum. This is a contradiction and thus, \mathbf{x}_s is a local minimum. So, no saddle points exist.

Theorem 1 implies that $L(\mathbf{x}^0)$ is nonempty and singleton.

⁴An alternative way to get a contradiction is to assume there exists a strict local maximum and use only convexity, instead of strict convexity.

2) We show that if we do not start at a local minimum, strict descent follows. Let $\mu(\cdot)$ denote the point-to-point mapping during one iteration step, i.e., $\mathbf{x}^{m+1} = \mu(\mathbf{x}^m)$. We show that if $\mathbf{x}^0 \notin C_J$, then $L(\mathbf{x}^0) \subseteq C_{J,\min}$. The result then follows by using the proof of the first part ⁵. To this end, let \mathbf{x}' be a fixed point under μ , i.e., $\mu(\mathbf{x}') = \mathbf{x}'$. Then, the subiteration steps of the algorithm yield $0 \in \partial_{\mathbf{x}_i} J_\lambda(\mathbf{x}'_1, \dots, \mathbf{x}'_k)$ for $i = 1, \dots, k$, which implies $0 \in \partial J_\lambda(\mathbf{x}')$, i.e., $\mathbf{x}' \in C_J$. The contrapositive implies that if $\mathbf{x} \notin C_J$, then $J_\lambda(\mu(\mathbf{x})) < J_\lambda(\mathbf{x})$ (strict descent). A simple induction on the number of iterations then concludes the proof. ■

APPENDIX H

LEMMA 4

The following technical lemma will be used in the proof of Lemma 5.

Lemma 4. *Let $\mathbf{z} \sim N(\mathbf{0}, \mathbf{A}_0 \otimes \mathbf{B}_0)$, where $\mathbf{A}_0 \in S_{++}^p$, $\mathbf{B}_0 \in S_{++}^f$. Then, for $m \geq 0$, we have the moment bound:*

$$\begin{aligned} & \mathbb{E} \left[\left(\sum_{i,j=1}^p \mathbf{X}_{i,j} ([\mathbf{z}]_{(i-1)f+k} [\mathbf{z}]_{(j-1)f+l} - [\mathbf{A}_0]_{i,j} [\mathbf{B}_0]_{k,l}) \right)^{m+2} \right] \\ & \leq (2m+2)!! p \left(\max_{1 \leq k \leq f} [\mathbf{B}_0]_{k,k} \|\mathbf{X}\|_2 \|\mathbf{A}_0\|_2 \right)^{m+2} \end{aligned}$$

Remark 4. *In the symmetric $\mathbf{X} \in S^p$ case, the bound in Lemma 4 can be tightened to*

$$\begin{aligned} & \mathbb{E} \left[\left(\sum_{i,j=1}^p \mathbf{X}_{i,j} ([\mathbf{z}]_{(i-1)f+k} [\mathbf{z}]_{(j-1)f+l} - [\mathbf{A}_0]_{i,j} [\mathbf{B}_0]_{k,l}) \right)^{m+2} \right] \\ & \leq (2m+2)!! \left(\max_{1 \leq k \leq f} [\mathbf{B}_0]_{k,k} \right)^{m+2} \text{tr}((\mathbf{X}\mathbf{A}_0)^{m+2}) \end{aligned}$$

Proof:

Consider the index set $\{\{i_1, j_1\}, \{i_2, j_2\}, \dots, \{i_{m+2}, j_{m+2}\}\}$. Define groups $G_k = \{i_k, j_k\}$ for $k = 1, \dots, m+2$. Let the generic notation $\pi(\cdot)$ denote the permutation operator of a set of indices.

Define the set of indices $M_{m+2} = M_{m+2}(i_1, j_1, \dots, i_{m+2}, j_{m+2})$ as the set containing sequences $(I_1, J_1, \dots, I_{m+2}, J_{m+2})$ satisfying the properties:

- 1) $\{I_1, J_1, \dots, I_{m+2}, J_{m+2}\}$ is a permutation of the index set $\{i_1, j_1, \dots, i_{m+2}, j_{m+2}\}$
 -i.e. $\{I_1, J_1, \dots, I_{m+2}, J_{m+2}\} = \pi(\{i_1, j_1, \dots, i_{m+2}, j_{m+2}\})$

⁵The first part of the proof showed $C_J = C_{J,\min}$.

- 2) For each $q \in \{1, \dots, m+2\}$, indices I_q and J_q must belong to disjoint groups $\{G_k\}_{k=1}^{m+2}$
- 3) Suppose a sequence $\{I_1, J_1, \dots, I_{m+2}, J_{m+2}\}$ satisfies the first two properties. Then, add it to M_{m+2} and M_{m+2} does not contain (block-permuted) sequences of the form
- $$\{\pi(\{\pi(\{I_1, J_1\}), \pi(\{I_2, J_2\}), \dots, \pi(\{I_{m+2}, J_{m+2}\})\})\}$$

It can be shown that $\text{card}(M_{m+2}) = (2m+2)!!$.

As an illustrative example, consider the case $m = 1$.

Example 1. For $m = 1$, the set M_{m+2} contains the following $4!! = 8$ elements:

$$\begin{aligned} & \{\{i_1, i_2\}, \{j_1, i_3\}, \{j_2, j_3\}\}, \{\{i_1, i_2\}, \{j_1, j_3\}, \{j_2, i_3\}\}, \\ & \{\{i_1, j_2\}, \{j_1, i_3\}, \{i_2, j_3\}\}, \{\{i_1, j_2\}, \{j_1, j_3\}, \{i_2, i_3\}\}, \\ & \{\{i_1, i_3\}, \{j_1, i_2\}, \{j_2, j_3\}\}, \{\{i_1, i_3\}, \{j_1, j_2\}, \{i_2, j_3\}\}, \\ & \{\{i_1, j_3\}, \{j_1, j_2\}, \{i_2, i_3\}\}, \{\{i_1, j_3\}, \{j_1, i_2\}, \{j_2, i_3\}\}. \end{aligned}$$

Of course, other equivalent possibilities for M_{m+2} are possible.

Note that $\text{tr}((\mathbf{X}\mathbf{A}_0)^{m+2}) \geq 0$ for all $m \geq 0$. From Isserlis' formula [23], we have:

$$\begin{aligned} & \mathbb{E} \left[\left(\sum_{i,j=1}^p \mathbf{X}_{i,j} ([\mathbf{z}]_{(i-1)f+k} [\mathbf{z}]_{(j-1)f+l} - [\mathbf{A}_0]_{i,j} [\mathbf{B}_0]_{k,l}) \right)^{m+2} \right] \\ &= \sum_{i_1, j_1=1}^p \cdots \sum_{i_{m+2}, j_{m+2}=1}^p \mathbf{X}_{i_1, j_1} \cdots \mathbf{X}_{i_{m+2}, j_{m+2}} \\ & \times \mathbb{E} \left[\prod_{\alpha=1}^{m+2} \left([\mathbf{z}]_{(i_\alpha-1)f+k} [\mathbf{z}]_{(j_\alpha-1)f+l} - [\mathbf{A}_0]_{i_\alpha, j_\alpha} [\mathbf{B}_0]_{k,l} \right) \right] \\ & \leq \left(\max_{1 \leq k \leq f} [\mathbf{B}_0]_{k,k} \right)^{m+2} \sum_{i_1, j_1=1}^p \cdots \sum_{i_{m+2}, j_{m+2}=1}^p \mathbf{X}_{i_1, j_1} \cdots \mathbf{X}_{i_{m+2}, j_{m+2}} \\ & \times \sum_{\{I_q, J_q\}_{q=1}^{m+2} \in M_{m+2}} \prod_{q=1}^{m+2} [\mathbf{A}_0]_{I_q, J_q} \\ & \leq \left(\max_{1 \leq k \leq f} [\mathbf{B}_0]_{k,k} \right)^{m+2} (2m+2)!! p (\|\mathbf{X}\|_2 \|\mathbf{A}_0\|_2)^{m+2} \end{aligned}$$

■

APPENDIX I

LEMMA 5

The following lemma will be used in the proof of Theorem 3, Theorem 4 and Theorem 5. The method of proof is by moment generating functions. A similar bound can be obtained under the same set of assumptions using standard decoupling arguments and Gaussian chaos Talagrand-based bounds.

Lemma 5. *Let \mathbf{X} be a $p \times p$ data-independent matrix. Define the linear operator \mathbf{T} as $\mathbf{T}(\mathbf{X}) = \hat{\mathbf{B}}(\mathbf{X}^{-1})$, where $\hat{\mathbf{B}}(\cdot)$ is defined in (8). Assume $\max_k [\mathbf{B}_0]_{k,k}$, $\|\mathbf{X}\|_2$, $\|\mathbf{A}_0\|_2$ are uniformly bounded constants as $p, f \rightarrow \infty$. Define $\mathbf{B}_* := \frac{\text{tr}(\mathbf{X}\mathbf{A}_0)}{p} \mathbf{B}_0$. Let $c, \tau > 0$. Define $\psi(u) = \sum_{m=0}^{\infty} \frac{(2m+2)!!}{m!} u^m$ ⁶. Let $\bar{C} := \frac{4(2+\tau)^2 \max(2,c)}{\psi(\frac{1}{2+\tau})} < \frac{np}{\log(\max(f,n))}$ ⁷. Then, with probability $1 - \frac{2}{\max(f,n)^c}$,*

$$|\mathbf{T}(\mathbf{X}) - \mathbf{B}_*|_{\infty} \leq \bar{k} \cdot \sqrt{4\psi\left(\frac{1}{2+\tau}\right) \max(2,c)} \sqrt{\frac{\log(\max(f,n))}{np}}$$

where $\bar{k} = \max_k [\mathbf{B}_0]_{k,k} \cdot \|\mathbf{X}\|_2 \|\mathbf{A}_0\|_2$.

Remark 5. *Choosing $c \leq 2$ in Lemma 5, the best relative constant is obtained by taking τ to infinity, which yields $\sqrt{4\psi(\frac{1}{2+\tau}) \max(2,c)} \rightarrow 4$.*

Remark 6. *For the case of symmetric matrices $\mathbf{X} \in S^p$, the constant \bar{k} can be improved to $\max_k [\mathbf{B}_0]_{k,k} \cdot \|\mathbf{X}\mathbf{A}_0\|_2$.*

Proof: This proof is based on a large-deviation theory argument. Fix $(k, l) \in \{1, \dots, f\}^2$. Note that $\mathbb{E}[\mathbf{T}(\mathbf{X})] = \mathbf{B}_*$. First we bound the upper tail probability on the difference $\mathbf{T}(\mathbf{X}) - \mathbf{B}_*$ and then we turn

⁶The double factorial notation is defined as

$$m!! = \begin{cases} m \cdot (m-2) \cdot \dots \cdot 3 \cdot 1 & \text{if } m > 0 \text{ is odd} \\ m \cdot (m-2) \cdot \dots \cdot 4 \cdot 2 & \text{if } m > 0 \text{ is even} \\ 1 & \text{if } m = -1 \text{ or } m = 0 \end{cases}$$

⁷If $p = f = n^{c'}$ for some $c' > 0$, this condition will hold for n large enough.

to the lower tail probability. Bounding the upper tail by using Markov's inequality, we have

$$\begin{aligned}
& \Pr([T(\mathbf{X})]_{k,l} - [\mathbf{B}_*]_{k,l} > \epsilon) \\
&= \Pr\left(\frac{1}{p} \sum_{i,j=1}^p \mathbf{X}_{i,j} [\hat{\mathbf{S}}_n(j,i)]_{k,l} - \frac{\text{tr}(\mathbf{X}\mathbf{A}_0)}{p} [\mathbf{B}_0]_{k,l} > \epsilon\right) \\
&= \Pr\left(\sum_{m=1}^n \sum_{i,j=1}^p \mathbf{X}_{i,j} \left([\mathbf{z}_m]_{(i-1)f+k} [\mathbf{z}_m]_{(j-1)f+l} \right. \right. \\
&\quad \left. \left. - [\mathbf{A}_0]_{i,j} [\mathbf{B}_0]_{k,l}\right) > np\epsilon\right) \\
&= \Pr\left(\exp\left\{t \sum_{m=1}^n \sum_{i,j=1}^p \mathbf{X}_{i,j} \left([\mathbf{z}_m]_{(i-1)f+k} [\mathbf{z}_m]_{(j-1)f+l} \right. \right. \right. \\
&\quad \left. \left. \left. - [\mathbf{A}_0]_{i,j} [\mathbf{B}_0]_{k,l}\right)\right\} > \exp\{tnp\epsilon\}\right) \\
&\leq e^{-tnp\epsilon} \mathbb{E}\left[\prod_{m=1}^n \exp\left\{t \sum_{i,j=1}^p \mathbf{X}_{i,j} \left([\mathbf{z}_m]_{(i-1)f+k} [\mathbf{z}_m]_{(j-1)f+l} \right. \right. \right. \\
&\quad \left. \left. \left. - [\mathbf{A}_0]_{i,j} [\mathbf{B}_0]_{k,l}\right)\right\}\right] \\
&\leq e^{-tnp\epsilon} \left(\mathbb{E}\left[\exp\left\{t\tilde{Y}^{(k,l)}\right\}\right]\right)^n \tag{28}
\end{aligned}$$

where we used the i.i.d. property of the data in (28) and $\tilde{Y}^{(k,l)} := \sum_{i,j=1}^p \mathbf{X}_{i,j} ([\mathbf{z}]_{(i-1)f+k} [\mathbf{z}]_{(j-1)f+l} - [\mathbf{A}_0]_{i,j} [\mathbf{B}_0]_{k,l})$. Define $p^2 \times 1$ random vector $\mathbf{z}^{(k,l)}$ as $[\mathbf{z}^{(k,l)}]_{(i-1)p+j} := [\mathbf{z}]_{(i-1)f+k} [\mathbf{z}]_{(j-1)f+l} - [\mathbf{A}_0]_{i,j} [\mathbf{B}_0]_{k,l}$ for $1 \leq i, j \leq p$. Clearly, this random vector is zero mean. The expectation term inside the parentheses in (28) is the MGF of the random variable $\tilde{Y}^{(k,l)} = \text{vec}(\mathbf{X})^T \mathbf{z}^{(k,l)}$. For notational simplicity, let $\tilde{\phi}_Y(t) = \mathbb{E}[e^{tY}]$ denote the MGF of a random vector Y . As a result, $\mathbb{E}[e^{t\tilde{Y}^{(k,l)}}] = \tilde{\phi}_{\tilde{Y}^{(k,l)}}(t)$.

Performing a second order Taylor expansion on $\tilde{\phi}_{\tilde{Y}^{(k,l)}}$ about the origin, we obtain:

$$\tilde{\phi}_{\tilde{Y}^{(k,l)}}(t) = \tilde{\phi}_{\tilde{Y}^{(k,l)}}(0) + \frac{d\tilde{\phi}_{\tilde{Y}^{(k,l)}}(0)}{dt} t + \frac{1}{2} \frac{d^2\tilde{\phi}_{\tilde{Y}^{(k,l)}}(\delta t)}{dt^2} t^2$$

for some $\delta \in [0, 1]$. Trivially, $\tilde{\phi}_{\tilde{Y}^{(k,l)}}(0) = 1$ and $\frac{d\tilde{\phi}_{\tilde{Y}^{(k,l)}}(0)}{dt} = \mathbb{E}[\text{vec}(\mathbf{X})^T \mathbf{z}^{(k,l)}] = 0$. Using the linearity of the expectation operator, we have:

$$\begin{aligned}
\frac{d^2\tilde{\phi}_{\tilde{Y}^{(k,l)}}(\delta t)}{dt^2} &= \mathbb{E}[(\tilde{Y}^{(k,l)})^2 e^{t\delta\tilde{Y}^{(k,l)}}] \\
&= \sum_{m=0}^{\infty} \frac{(\delta t)^m}{m!} \mathbb{E}[(\text{vec}(\mathbf{X})^T \mathbf{z}^{(k,l)})^{m+2}]
\end{aligned}$$

Using the elementary inequality $1 + y \leq e^y$ for $y > -1$, and after some algebra, we have:

$$n \ln(\tilde{\phi}_{\tilde{Y}^{(k,l)}}(t)) \leq \frac{n}{2} t^2 \sum_{m=0}^{\infty} T_m(t) \tag{29}$$

where $T_m(t) := \frac{(t\delta)^m}{m!} \mathbb{E}[(\text{vec}(\mathbf{X})^T \mathbf{z}^{(k,l)})^{m+2}]$. Note that

$$\begin{aligned}
t^2 T_m(t) &\leq \frac{t^{m+2}}{m!} \mathbb{E} \left[\left(\sum_{i,j=1}^p \mathbf{X}_{i,j} ([\mathbf{z}]_{(i-1)f+k} [\mathbf{z}]_{(j-1)f+l} \right. \right. \\
&\quad \left. \left. - [\mathbf{A}_0]_{i,j} [\mathbf{B}_0]_{k,l} \right)^{m+2} \right] \\
&= \frac{t^{m+2}}{m!} \sum_{i_1, j_1=1}^p \cdots \sum_{i_{m+2}, j_{m+2}=1}^p \mathbf{X}_{i_1, j_1} \cdots \mathbf{X}_{i_{m+2}, j_{m+2}} \\
&\quad \times \mathbb{E} \left[\prod_{\alpha=1}^{m+2} \left([\mathbf{z}]_{(i_\alpha-1)f+k} [\mathbf{z}]_{(j_\alpha-1)f+l} - [\mathbf{A}_0]_{i_\alpha, j_\alpha} [\mathbf{B}_0]_{k,l} \right) \right] \\
&\leq \frac{t^{m+2}}{m!} (2m+2)!! p \left(\max_{1 \leq k \leq f} [\mathbf{B}_0]_{k,k} \|\mathbf{X}\|_2 \|\mathbf{A}_0\|_2 \right)^{m+2} \\
&= \frac{(2m+2)!!}{m!} (t\bar{k})^{m+2} p
\end{aligned} \tag{30}$$

where (30) follows from Lemma 4⁸. Also, we defined $\bar{k} = \max_{1 \leq k \leq f} [\mathbf{B}_0]_{k,k} \cdot \|\mathbf{X}\|_2 \|\mathbf{A}_0\|_2$. Summing the result over m , and letting $u := t\bar{k} > 0$, $a_m(u) := \frac{(2m+2)!!}{m!} u^m$, $\psi(u) := \sum_{m=0}^{\infty} a_m(u)$, we obtain:

$$t^2 \sum_{m=0}^{\infty} T_m(t) \leq pu^2 \psi(u) \Big|_{u=t\bar{k}} \tag{31}$$

By the ratio test [21], the infinite series $\sum_{m=0}^{\infty} a_m(u)$ converges if $u < 1/2$. To see this, note

$$\begin{aligned}
\rho &:= \lim_{m \rightarrow \infty} \frac{a_{m+1}(u)}{a_m(u)} \\
&= \lim_{m \rightarrow \infty} u \frac{(2m+4)!!}{(2m+2)!!} \frac{m!}{(m+1)!} \\
&= \lim_{m \rightarrow \infty} 2u \frac{1+2/m}{1+1/m} \\
&= 2u < 1
\end{aligned}$$

Using (31) in (29), and the result in (28), we obtain the exponential bound:

$$\begin{aligned}
&\Pr([\mathbf{T}(\mathbf{X})]_{k,l} - [\mathbf{B}_*]_{k,l} > \epsilon) \\
&\leq \exp \left\{ -tnp\epsilon + \frac{np(t\bar{k})^2}{2} \psi(t\bar{k}) \right\}
\end{aligned}$$

Let $t < \frac{1}{(2+\tau)\bar{k}}$ and $\epsilon < \frac{1}{2+\tau} \psi(\frac{1}{2+\tau})\bar{k} < \infty$. By the monotonicity of $\psi(\cdot)$, we have:

$$\Pr([\mathbf{T}(\mathbf{X})]_{k,l} - [\mathbf{B}_*]_{k,l} > \epsilon) \leq \exp \left\{ -tnp\epsilon + \frac{npt^2\bar{k}^2}{2} \psi\left(\frac{1}{2+\tau}\right) \right\} \tag{32}$$

⁸In the symmetric \mathbf{X} case, this bound can be tightened using $\text{tr}((\mathbf{X}\mathbf{A}_0)^{m+2}) \leq p(\|\mathbf{X}\mathbf{A}_0\|_2)^{m+2}$.

Optimizing (32) over t , we obtain $t^* = \frac{\epsilon}{\bar{k}^2 \psi(\frac{1}{2+\tau})}$. Clearly, $t^* < \frac{1}{(2+\tau)\bar{k}}$. Plugging this into (32), we obtain:

$$\Pr([\mathbf{T}(\mathbf{X})]_{k,l} - [\mathbf{B}_*]_{k,l} > \epsilon) \leq \exp \left\{ - \frac{np\epsilon^2}{2\bar{k}^2 \psi(\frac{1}{2+\tau})} \right\}$$

Define $C := \frac{1}{2\bar{k}^2 \psi(\frac{1}{2+\tau})}$. Since $\psi(\frac{1}{2+\tau}) < \infty$, $C > 0$. Thus, for all $\epsilon < \frac{1}{2+\tau} \psi(\frac{1}{2+\tau}) \bar{k}$, we have

$$P([\mathbf{T}(\mathbf{X})]_{k,l} - [\mathbf{B}_*]_{k,l} > \epsilon) \leq e^{-np\epsilon^2 C} \quad (33)$$

where $C > 0$ is independent of n, p, f .

Next, we bound the lower tail:

$$\begin{aligned} & \Pr([\mathbf{T}(\mathbf{X})]_{k,l} - \mathbb{E}[\mathbf{T}(\mathbf{X})]_{k,l} < -\epsilon) \\ &= \Pr \left(\sum_{m=1}^n \sum_{i,j=1}^p -\mathbf{X}_{i,j}([\mathbf{z}_m]_{(j-1)f+k}[\mathbf{z}_m]_{(i-1)f+l} \right. \\ & \quad \left. - [\mathbf{A}_0]_{i,j}[\mathbf{B}_0]_{k,l}) > np\epsilon \right) \\ &\leq e^{-tnp\epsilon} \left(\tilde{\phi}_{\tilde{Y}^{(k,l)}}(-t) \right)^n \end{aligned}$$

where $\tilde{\phi}_{\tilde{Y}^{(k,l)}}$ is the MGF of $\tilde{Y}^{(k,l)}$. Performing a second order Taylor expansion as before, we have:

$$\begin{aligned} \tilde{\phi}_{\tilde{Y}^{(k,l)}}(-t) &= \tilde{\phi}_{\tilde{Y}^{(k,l)}}(0) - \frac{d\tilde{\phi}_{\tilde{Y}^{(k,l)}}(0)}{dt}t + \frac{1}{2} \frac{d^2\tilde{\phi}_{\tilde{Y}^{(k,l)}}(0)}{dt^2}t^2 \\ &= 1 + \frac{t^2}{2} \sum_{m=0}^{\infty} T'_m(t) \end{aligned}$$

where $T'_m(t) := \frac{(-t\delta)^m}{m!} \mathbb{E}[(\langle \text{vec}(\mathbf{X}), \mathbf{z}^{(k,l)} \rangle)^{m+2}] = (-1)^m T_m(t) \leq T_m(t)$ and $\delta \in [0, 1]$. Proceeding similarly as above, it can be shown that for all $\epsilon < \frac{1}{2+\tau} \psi(\frac{1}{2+\tau}) \bar{k}$:

$$\Pr([\mathbf{T}(\mathbf{X})]_{k,l} - \mathbb{E}[\mathbf{T}(\mathbf{X})]_{k,l} < -\epsilon) \leq e^{-np\epsilon^2 C} \quad (34)$$

where C was defined as before. From (33) and (34), we conclude that for all $\epsilon < \frac{1}{2+\tau} \psi(\frac{1}{2+\tau}) \bar{k}$:

$$\begin{aligned} & \Pr(|[\mathbf{T}(\mathbf{X})]_{k,l} - \mathbb{E}[\mathbf{T}(\mathbf{X})]_{k,l}| > \epsilon) \\ &\leq \Pr([\mathbf{T}(\mathbf{X})]_{k,l} - \mathbb{E}[\mathbf{T}(\mathbf{X})]_{k,l} > \epsilon) \\ & \quad + \Pr([\mathbf{T}(\mathbf{X})]_{k,l} - \mathbb{E}[\mathbf{T}(\mathbf{X})]_{k,l} < -\epsilon) \\ &\leq 2e^{-np\epsilon^2 C} \end{aligned}$$

The union bound over $(k, l) \in \{1, \dots, f\}^2$ completes the proof. Let us rewrite this. If $\frac{4 \max(2, c) \log(\max(f, n))(2+\tau)^2}{\psi(\frac{1}{2+\tau})} < np$, then with probability $1 - \frac{2}{\max(f, n)^c}$,

$$\begin{aligned} \|\mathbf{T}(\mathbf{X}) - \mathbb{E}[\mathbf{T}(\mathbf{X})]\|_\infty &\leq \bar{k} \cdot \sqrt{2\psi\left(\frac{1}{2+\tau}\right)} \sqrt{\frac{\log((2f^2)/(2/\max(f, n)^c))}{np}} \\ &\leq \bar{k} \cdot 2 \sqrt{\psi\left(\frac{1}{2+\tau}\right) \max(2, c)} \sqrt{\frac{\log \max(f, n)}{np}} \end{aligned}$$

■

APPENDIX J

PROPOSITION 1

Proposition 1. Let $\mathbf{S}_{p,f,n}$ be a $d' \times d'$ (where $d' = p$ or $d' = f$) random matrix such that with probability $1 - \frac{2}{n^2}$, $\|\mathbf{S}_{p,f,n} - \boldsymbol{\Sigma}_*\|_\infty \leq r_{p,f,n}$. Assume $\boldsymbol{\Sigma}_* \in S_{++}^{d'}$ has uniformly bounded spectrum as $p, f \rightarrow \infty$ (analog to Assumption 1). Choose $\lambda_{p,f,n} = c \cdot r_{p,f,n}$ for some absolute constant $c > 0$. Consider the Glasso operator $\mathbf{G}(\cdot, \cdot)$ defined in (10). Let $s = s_{\boldsymbol{\Theta}_*}$ be the sparsity parameter associated with $\boldsymbol{\Theta}_* := \boldsymbol{\Sigma}_*^{-1}$. Assume $\sqrt{d' + s} \cdot r_{p,f,n} = o(1)$. Then, with probability $1 - \frac{2}{n^2}$,

$$\|\mathbf{G}(\mathbf{S}_{p,f,n}, \lambda_{p,f,n}) - \boldsymbol{\Theta}_*\|_F \leq \frac{2\sqrt{2}(1+c)}{\lambda_{\min}(\boldsymbol{\Sigma}_*)^2} \sqrt{d' + s} \cdot r_{p,f,n}$$

as $n \rightarrow \infty$.

Proof: The proof follows from a slight modification of Thm. 1 in [17], or Thm. 3 in [16]. This modification is due to the different $r_{p,f,n}$. ■

APPENDIX K

PROOF OF THEOREM 3

Proof: As in the proof of Thm. 1 in [5], let $\mathbf{B}_* = \frac{\text{tr}(\mathbf{A}_0 \mathbf{A}_{init}^{-1})}{p} \mathbf{B}_0$ and $\mathbf{A}_* = (\frac{\text{tr}(\mathbf{A}_0 \mathbf{A}_{init}^{-1})}{p})^{-1} \mathbf{A}_0$. Note that Assumption 1 implies that $\|\mathbf{B}_*\|_2 = \Theta(1)$ and $\|\mathbf{A}_*\|_2 = \Theta(1)$ as $p, f \rightarrow \infty$.

For concreteness, we first present the result for $k = 3$ iterations. Then, we generalize the analysis to $k \geq 3$ flip-flop iterations.

Define intermediate error matrices:

$$\begin{aligned} \tilde{\mathbf{B}}^0 &= \hat{\mathbf{B}}(\mathbf{A}_{init}) - \mathbf{B}_* \\ \tilde{\mathbf{A}}^1 &= \hat{\mathbf{A}}(\hat{\mathbf{B}}(\mathbf{A}_{init})) - \mathbf{A}_* \\ \tilde{\mathbf{B}}^2 &= \hat{\mathbf{B}}(\hat{\mathbf{A}}(\hat{\mathbf{B}}(\mathbf{A}_{init}))) - \mathbf{B}_* \end{aligned}$$

Lemma 5 implies $|\tilde{\mathbf{B}}^0|_\infty = O_P\left(\sqrt{\frac{\log M}{np}}\right)$ as $n, p, f \rightarrow \infty$, so we have

$$\|\tilde{\mathbf{B}}^0\|_F = O_P\left(\sqrt{\frac{f^2 \log M}{np}}\right) \quad (35)$$

From the matrix inversion lemma [24], we have

$$\begin{aligned} \hat{\mathbf{B}}(\mathbf{A}_{init})^{-1} &= (\mathbf{B}_* + \tilde{\mathbf{B}}^0)^{-1} \\ &= \mathbf{B}_*^{-1} - \mathbf{B}_*^{-1} \tilde{\mathbf{B}}^0 \mathbf{B}_*^{-1} + \Delta_1 \end{aligned} \quad (36)$$

where

$$\Delta_1 = \sum_{k=2}^{\infty} (-\mathbf{B}_*^{-1} \tilde{\mathbf{B}}^0)^k \mathbf{B}_*^{-1}$$

Let us show

$$\|\Delta_1\|_F = O_P\left(\frac{f^{2.5} \log M}{np}\right) \quad (37)$$

From the submultiplicative property of the spectral norm, the triangle inequality and (35), we have:

$$\begin{aligned} \|\Delta_1\|_2 &= \left\| \sum_{k=2}^{\infty} (-\mathbf{B}_*^{-1} \tilde{\mathbf{B}}^0)^k \mathbf{B}_*^{-1} \right\|_2 \\ &\leq \sum_{k=2}^{\infty} (\|\mathbf{B}_*^{-1}\|_2 \|\tilde{\mathbf{B}}^0\|_2)^k \|\mathbf{B}_*^{-1}\|_2 \\ &\leq \sum_{k=2}^{\infty} (\|\mathbf{B}_*^{-1}\|_2 \|\tilde{\mathbf{B}}^0\|_F)^k \|\mathbf{B}_*^{-1}\|_2 \\ &= O_P\left(\frac{f^2 \log f}{np}\right) \end{aligned}$$

where we used Assumption 1 and $b_{p,f,n} = \|\mathbf{B}_*^{-1}\|_2 \|\tilde{\mathbf{B}}^0\|_F = O_P\left(\sqrt{\frac{f^2 \log M}{np}}\right)$ in the bound:

$$\sum_{k=2}^{\infty} b_{p,f,n}^k = b_{p,f,n}^2 + \frac{b_{p,f,n}^3}{1 - b_{p,f,n}} \leq C b_{p,f,n}^2$$

with high probability as $p, f, n \rightarrow \infty$ (here, $C > 0$ is an absolute constant). This type of asymptotic argument will be used in the remainder of the proof to retain the dominant terms. The inequality $\|\Delta_1\|_F \leq \sqrt{f} \|\Delta_1\|_2$ then implies (37).

Next we expand $\text{vec}(\tilde{\mathbf{A}}^1)$:

$$\begin{aligned}
\text{vec}(\tilde{\mathbf{A}}^1) &= \frac{1}{f} \hat{\mathbf{R}}_A \text{vec}(\hat{\mathbf{B}}(\mathbf{A}_{init})^{-1}) - \text{vec}(\mathbf{A}_*) \\
&= \frac{1}{f} \tilde{\mathbf{R}}_A \text{vec}(\mathbf{B}_*^{-1}) - \frac{1}{f} \mathbf{R}_A (\mathbf{B}_*^{-T} \otimes \mathbf{B}_*^{-1}) \text{vec}(\tilde{\mathbf{B}}^0) \\
&\quad + \text{vec}(\mathbf{\Delta}_2) \\
&= \text{vec}(\hat{\mathbf{A}}(\mathbf{B}_*) - \mathbf{A}_*) - \frac{\text{tr}(\mathbf{B}_*^{-1} \tilde{\mathbf{B}}^0)}{f} \text{vec}(\mathbf{A}_*) \\
&\quad + \text{vec}(\mathbf{\Delta}_2)
\end{aligned} \tag{38}$$

where we used (see (106) from [5])

$$\begin{aligned}
\mathbf{R}_A (\mathbf{B}_*^{-T} \otimes \mathbf{B}_*^{-1}) &= \text{vec}(\mathbf{A}_*) \text{vec}(\mathbf{B}_*^{-T})^T \\
\mathbf{R}_B (\mathbf{A}_*^{-T} \otimes \mathbf{A}_*^{-1}) &= \text{vec}(\mathbf{B}_*) \text{vec}(\mathbf{A}_*^{-T})^T
\end{aligned} \tag{39}$$

and $\mathbf{\Delta}_2$ is given by

$$\begin{aligned}
\text{vec}(\mathbf{\Delta}_2) &= \frac{1}{f} \hat{\mathbf{R}}_A \text{vec}(\mathbf{\Delta}_1) \\
&\quad - \frac{1}{f} \tilde{\mathbf{R}}_A (\mathbf{B}_*^{-T} \otimes \mathbf{B}_*^{-1}) \text{vec}(\tilde{\mathbf{B}}^0)
\end{aligned} \tag{40}$$

Using (37) and (35) and bounds $\|\tilde{\mathbf{R}}_A\|_2 \leq pf|\hat{\mathbf{S}}_n - \mathbf{\Sigma}_0|_\infty = O_P\left(pf\sqrt{\frac{\log(pf)}{n}}\right)$ (see Lemma 1 in [17]) and $\|\mathbf{R}_A\|_2 = \|\mathbf{A}_0\|_F \|\mathbf{B}_0\|_F \leq \sqrt{pf} \|\mathbf{\Sigma}_0\|_2$:

$$\begin{aligned}
\|\mathbf{\Delta}_2\|_F &\leq \frac{1}{f} \|\hat{\mathbf{R}}_A\|_2 \|\mathbf{\Delta}_1\|_F + \frac{1}{f} \|\tilde{\mathbf{R}}_A\|_2 \|\mathbf{B}_*^{-1}\|_2^2 \|\tilde{\mathbf{B}}^0\|_F \\
&= O_P\left((f^2 p^{-1/2} + \sqrt{p}) \frac{\log M}{n}\right)
\end{aligned} \tag{41}$$

Now, using the triangle inequality in (38) and the bounds (41) and (35), we obtain:

$$\begin{aligned}
\|\tilde{\mathbf{A}}^1\|_F &\leq \|\hat{\mathbf{A}}(\mathbf{B}_*) - \mathbf{A}_*\|_F + \frac{\|\mathbf{A}_*\|_F \|\mathbf{B}_*^{-1}\|_F \|\tilde{\mathbf{B}}^0\|_F}{f} + \|\mathbf{\Delta}_2\|_F \\
&\leq O_P\left(\underbrace{(pf^{-1/2} + \sqrt{f}) \sqrt{\frac{\log M}{n}}}_{a_{p,f,n}^{(1)}} + \underbrace{(f^2 p^{-1/2} + \sqrt{p}) \frac{\log M}{n}}_{\tilde{a}_{p,f,n}^{(1)}}\right)
\end{aligned} \tag{42}$$

where we also used the Cauchy-Schwarz inequality $|\text{tr}(\mathbf{A}_* \mathbf{B}_*^{-1})| \leq \|\mathbf{A}_*\|_F \|\mathbf{B}_*^{-1}\|_F \leq \sqrt{pf} \|\mathbf{A}_*\|_2 \|\mathbf{B}_*^{-1}\|_2$ and Lemma 5. The growth assumption implies $\frac{\tilde{a}_{p,f,n}^{(1)}}{a_{p,f,n}^{(1)}} = \frac{f^2 p^{-1/2} + \sqrt{p}}{pf^{-1/2} + \sqrt{f}} \sqrt{\frac{\log M}{n}} = O(1)$. This implies that $\tilde{a}_{p,f,n}^{(1)}$ is asymptotically negligible and thus:

$$\|\tilde{\mathbf{A}}^1\|_F = O_P\left((pf^{-1/2} + \sqrt{f}) \sqrt{\frac{\log M}{n}}\right) \tag{43}$$

In parallel to (36), we expand:

$$\begin{aligned} (\hat{\mathbf{A}}(\hat{\mathbf{B}}(\mathbf{A}_{init})))^{-1} &= (\mathbf{A}_* + \tilde{\mathbf{A}}^1)^{-1} \\ &= \mathbf{A}_*^{-1} - \mathbf{A}_*^{-1} \tilde{\mathbf{A}}^1 \mathbf{A}_*^{-1} + \Delta_3 \end{aligned} \quad (44)$$

where

$$\Delta_3 = \sum_{k=2}^{\infty} (-\mathbf{A}_*^{-1} \tilde{\mathbf{A}}^1)^k \mathbf{A}_*^{-1} \quad (45)$$

Then, expanding $\tilde{\mathbf{B}}^2$ in parallel to (38),

$$\begin{aligned} \text{vec}(\tilde{\mathbf{B}}^2) &= \frac{1}{p} \hat{\mathbf{R}}_B \text{vec}((\hat{\mathbf{A}}(\hat{\mathbf{B}}(\mathbf{A}_{init})))^{-1}) - \text{vec}(\mathbf{B}_*) \\ &= \text{vec}(\hat{\mathbf{B}}(\mathbf{A}_*) - \mathbf{B}_*) - \frac{\text{tr}(\mathbf{A}_*^{-1} \tilde{\mathbf{A}}^1)}{p} \text{vec}(\mathbf{B}_*) \\ &\quad + \text{vec}(\Delta_4) \end{aligned} \quad (46)$$

where

$$\text{vec}(\Delta_4) = \frac{1}{p} \hat{\mathbf{R}}_B \text{vec}(\Delta_3) - \frac{1}{p} \tilde{\mathbf{R}}_B (\mathbf{A}_*^{-T} \otimes \mathbf{A}_*^{-1}) \text{vec}(\tilde{\mathbf{A}}^1) \quad (47)$$

Using similar techniques as above, it follows from (43) and (45):

$$\|\Delta_3\|_2 = O_P \left((pf^{-1/2} + \sqrt{f})^2 \frac{\log M}{n} \right)$$

which implies

$$\|\Delta_3\|_F = O_P \left(\sqrt{p} (pf^{-1/2} + \sqrt{f})^2 \frac{\log M}{n} \right) \quad (48)$$

Applying the triangle inequality in (47) and using (43), (48), we have after some algebra:

$$\begin{aligned} \|\Delta_4\|_F &\leq \frac{1}{p} \|\hat{\mathbf{R}}_B\|_2 \|\Delta_3\|_F + \frac{1}{p} \|\tilde{\mathbf{R}}_B\|_2 \|\mathbf{A}_*^{-1}\|_2^2 \|\tilde{\mathbf{A}}^1\|_F \\ &= O_P \left(\sqrt{f} (pf^{-1/2} + \sqrt{f})^2 \frac{\log M}{n} \right) \end{aligned} \quad (49)$$

From (46), (43) and (47), we obtain:

$$\begin{aligned} \|\tilde{\mathbf{B}}^2\|_F &\leq \|\hat{\mathbf{B}}(\mathbf{A}_*) - \mathbf{B}_*\|_F + \frac{|\text{tr}(\mathbf{A}_*^{-1} \tilde{\mathbf{A}}^1)|}{p} \|\mathbf{B}_*\|_F + \|\Delta_4\|_F \\ &\leq \|\hat{\mathbf{B}}(\mathbf{A}_*) - \mathbf{B}_*\|_F + \frac{\|\tilde{\mathbf{A}}^1\|_F \sqrt{pf}}{p} \|\mathbf{A}_*^{-1}\|_2 \|\mathbf{B}_*\|_2 + \|\Delta_4\|_F \\ &= O_P \left(\underbrace{(fp^{-1/2} + \sqrt{p}) \sqrt{\frac{\log M}{n}}}_{\tilde{b}^{(2)}} + \underbrace{\sqrt{f} (pf^{-1/2} + \sqrt{f})^2 \frac{\log M}{n}}_{\tilde{b}^{(2)}} \right) \\ &= O_P \left((fp^{-1/2} + \sqrt{p}) \sqrt{\frac{\log M}{n}} \right) \end{aligned} \quad (50)$$

where we used the growth assumption to obtain $\frac{\tilde{b}^{(2)}}{b^{(2)}} = \frac{\sqrt{f}(pf^{-1/2} + \sqrt{f})^2}{fp^{-1/2} + \sqrt{p}} = O(1)$.

Define the errors $\tilde{\mathbf{R}}_{FF}(3) = \hat{\mathbf{R}}_{FF}(3) - \Sigma_0$ and $\tilde{\mathbf{R}} := \hat{\mathbf{S}}_n - \Sigma_0$. Using the permutation operator \mathcal{R} defined in [5], we plug in (38) and (46) and simplify:

$$\begin{aligned} \text{vec}(\mathcal{R}(\tilde{\mathbf{R}}_{FF}(3))) &= \text{vec}(\text{vec}(\tilde{\mathbf{A}}^1)\text{vec}(\mathbf{B}_*)^T) \\ &\quad + \text{vec}(\text{vec}(\mathbf{A}_*)\text{vec}(\tilde{\mathbf{B}}^2)^T) + \text{vec}(\text{vec}(\tilde{\mathbf{A}}^1)\text{vec}(\tilde{\mathbf{B}}^2)^T) \\ &= \mathbf{P}_R^{-1}\Xi\text{vec}(\tilde{\mathbf{R}}) + \text{vec}(\tilde{\Delta}_R) \end{aligned} \quad (51)$$

where $\tilde{\Delta}_R$ is given by:

$$\begin{aligned} \tilde{\Delta}_R &= \text{vec}(\Delta_2)\text{vec}(\mathbf{B}_*)^T + \text{vec}(\mathbf{A}_*)\text{vec}(\Delta_4)^T \\ &\quad - \frac{\text{tr}(\Delta_2\mathbf{A}_*^{-1})}{p}\text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{B}_0)^T + \text{vec}(\tilde{\mathbf{A}}^1)\text{vec}(\tilde{\mathbf{B}}^2)^T \end{aligned} \quad (52)$$

and Ξ is a data-independent matrix:

$$\begin{aligned} \Xi &= \mathbf{P}_R \left[\frac{1}{f}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{B}_0^{-1})^T \otimes \mathbf{I}_{p^2})\mathbf{P}_{R_A} \right. \\ &\quad + \frac{1}{p}(\mathbf{I}_{f^2} \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{A}_0^{-1})^T)\mathbf{K}_{f^2, p^2}\mathbf{P}_{R_B} \\ &\quad - \frac{1}{pf}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{A}_0^{-1})^T \\ &\quad \left. \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{B}_0^{-1})^T)\mathbf{K}_{p^2, f^2}\mathbf{P}_{R_A} \right] \end{aligned} \quad (53)$$

From (51) and (24) in [5], we have:

$$\text{vec}(\tilde{\mathbf{R}}_{FF}(3)) = \Xi\text{vec}(\tilde{\mathbf{R}}) + \text{vec}(\Delta_R) \quad (54)$$

where $\text{vec}(\Delta_R) := \mathbf{P}_R\text{vec}(\tilde{\Delta}_R)$. Since \mathbf{P}_R is a permutation matrix, $\|\Delta_R\|_F = \|\tilde{\Delta}_R\|_F$. Using this and the triangle inequality in (52):

$$\begin{aligned} \|\Delta_R\|_F &\leq \|\Delta_2\|_F\sqrt{f}\|\mathbf{B}_*\|_2 + \sqrt{p}\|\mathbf{A}_*\|_2\|\Delta_4\|_F \\ &\quad + \|\Delta_2\|_F\|\mathbf{A}_*^{-1}\|_2\|\mathbf{A}_0\|_2\sqrt{f}\|\mathbf{B}_0\|_2 + \|\tilde{\mathbf{A}}^1\|_F\|\tilde{\mathbf{B}}^2\|_F \\ &= O_P\left(\sqrt{pf}\bar{M}\frac{\log M}{n}\right) \end{aligned} \quad (55)$$

where we used the Frobenius norm bounds in (41), (49), (42) and (50). Recall that \bar{M} was defined in the statement of the Theorem.

From (54), we obtain:

$$\begin{aligned} \text{Cov}(\text{vec}(\tilde{\mathbf{R}}_{FF}(3))) &= \Xi\text{Cov}(\text{vec}(\tilde{\mathbf{R}}))\Xi^T + \text{E}[\Xi\text{vec}(\tilde{\mathbf{R}})\text{vec}(\Delta_R)^T] \\ &\quad + \text{E}[\text{vec}(\Delta_R)(\Xi\text{vec}(\tilde{\mathbf{R}}))^T] + \text{E}[\text{vec}(\Delta_R)\text{vec}(\Delta_R)^T] \end{aligned} \quad (56)$$

where we used $E[\text{vec}(\tilde{\mathbf{R}})] = \mathbf{0}$ and $E[\text{vec}(\mathbf{\Delta}_R)] = \mathbf{0}$. By Jensen's inequality and the Cauchy-Schwarz inequality, we obtain:

$$\begin{aligned} |\text{tr}(E[\mathbf{\Xi} \text{vec}(\tilde{\mathbf{R}}) \text{vec}(\mathbf{\Delta}_R)^T])| &\leq E| \langle \text{vec}(\mathbf{\Delta}_R), \mathbf{\Xi} \text{vec}(\tilde{\mathbf{R}}) \rangle | \\ &\leq E\|\mathbf{\Xi} \text{vec}(\tilde{\mathbf{R}})\|_2 \|\text{vec}(\mathbf{\Delta}_R)\|_2 \\ &\leq \|\mathbf{\Xi}\|_2 E[\|\tilde{\mathbf{R}}\|_F \|\mathbf{\Delta}_R\|_F] \end{aligned} \quad (57)$$

From (53), it is easy to see that $\|\mathbf{\Xi}\|_2 = O(1)$ as $p, f \rightarrow \infty$. Similarly, we have

$$|\text{tr}(E[\text{vec}(\mathbf{\Delta}_R) \text{vec}(\mathbf{\Delta}_R)^T])| \leq E[\|\mathbf{\Delta}_R\|_F^2] \quad (58)$$

Next, we show

$$\begin{aligned} \text{tr}(\mathbf{\Xi}(\mathbf{\Sigma}_0 \otimes \mathbf{\Sigma}_0)\mathbf{\Xi}^T) &= \text{tr}(\mathbf{A}_0)^2 \frac{\text{tr}(\mathbf{B}_0^2)}{f} + \text{tr}(\mathbf{B}_0)^2 \frac{\text{tr}(\mathbf{A}_0^2)}{p} \\ &\quad - \frac{\text{tr}(\mathbf{A}_0^2) \text{tr}(\mathbf{B}_0^2)}{pf} \end{aligned} \quad (59)$$

Once (59) is established, using $\text{tr}(\mathbf{A}_0)^2 \leq \|\mathbf{A}_0\|_2^2 p^2$ and $\text{tr}(\mathbf{B}_0^2) \leq f \|\mathbf{B}_0\|_2^2$, we have:

$$\text{tr}(\mathbf{\Xi}(\mathbf{\Sigma}_0 \otimes \mathbf{\Sigma}_0)\mathbf{\Xi}^T) \leq (p^2 + f^2) \|\mathbf{\Sigma}_0\|_2^2 \quad (60)$$

Similarly, we have

$$\text{tr}(\mathbf{\Xi}(\mathbf{\Sigma}_0 \otimes \mathbf{\Sigma}_0)\mathbf{\Xi}^T) \geq (p^2 + f^2) \lambda_{\min}(\mathbf{A}_0)^2 \lambda_{\min}(\mathbf{B}_0)^2 - \|\mathbf{\Sigma}_0\|_2^2 \quad (61)$$

Let us prove (59) now. Note that

$$\text{tr}(\mathbf{\Xi}(\mathbf{\Sigma}_0 \otimes \mathbf{\Sigma}_0)\mathbf{\Xi}^T) = \text{vec}(\mathbf{\Xi}^T \mathbf{\Xi})^T \text{vec}(\mathbf{\Sigma}_0 \otimes \mathbf{\Sigma}_0) = \sum_{i=1}^9 T_i$$

where T_i will be defined and calculated below.

We have

$$\begin{aligned}
T_1 &= \text{vec}(\{\frac{1}{f}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{B}_0^{-1})^T \otimes \mathbf{I}_{p^2})\mathbf{P}_{R_A}\}^T \times \\
&\quad \{\frac{1}{f}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{B}_0^{-1})^T \otimes \mathbf{I}_{p^2})\mathbf{P}_{R_A}\})^T \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0) \\
&= \frac{\text{tr}(\mathbf{B}_0)^2}{f^2} \text{vec}((\text{vec}(\mathbf{B}_0^{-1})\text{vec}(\mathbf{B}_0^{-1})^T \otimes \mathbf{I}_{p^2}))^T \times \\
&\quad \text{vec}(\mathbf{P}_{R_A}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0)\mathbf{P}_{R_A}^T) \\
&= \frac{\text{tr}(\mathbf{B}_0)^2}{f^2} \text{vec}((\text{vec}(\mathbf{B}_0^{-1})\text{vec}(\mathbf{B}_0^{-1})^T \otimes \mathbf{I}_{p^2}))^T \times \\
&\quad \text{vec}((\mathbf{B}_0 \otimes \mathbf{B}_0) \otimes (\mathbf{A}_0 \otimes \mathbf{A}_0)) \\
&= \text{tr}(\mathbf{A}_0)^2 \frac{\text{tr}(\mathbf{B}_0^2)}{f^2} \cdot \text{vec}(\mathbf{B}_0^{-1})^T (\mathbf{B}_0 \otimes \mathbf{B}_0) \text{vec}(\mathbf{B}_0^{-1}) \\
&= \text{tr}(\mathbf{A}_0)^2 \frac{\text{tr}(\mathbf{B}_0^2)}{f}
\end{aligned}$$

where we used several Kronecker product identities [24] and the definition of the permutation matrix \mathbf{P}_{R_A} .

The rest of the terms are given by:

$$\begin{aligned}
T_2 &= \text{vec}(\{\frac{1}{f}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{Y}_0)^T \otimes \mathbf{I}_{p^2})\mathbf{P}_{R_A}\}^T \\
&\quad \times \{\frac{1}{p}(\mathbf{I}_{f^2} \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{X}_0)^T)\mathbf{K}_{f^2,p^2}\mathbf{P}_{R_B}\}^T \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0) \\
T_3 &= \text{vec}(\{\frac{1}{f}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{Y}_0)^T \otimes \mathbf{I}_{p^2})\mathbf{P}_{R_A}\}^T \\
&\quad \times \{\frac{-1}{pf}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{X}_0)^T \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{Y}_0)^T)\mathbf{K}_{p^2,f^2}\mathbf{P}_{R_A}\}^T \\
&\quad \times \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0) \\
T_4 &= \text{vec}(\{\frac{1}{p}(\mathbf{I}_{f^2} \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{X}_0)^T)\mathbf{K}_{f^2,p^2}\mathbf{P}_{R_B}\}^T \\
&\quad \times \{\frac{1}{f}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{Y}_0)^T \otimes \mathbf{I}_{p^2})\mathbf{P}_{R_A}\}^T \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0) \\
T_5 &= \text{vec}(\{\frac{1}{p}(\mathbf{I}_{f^2} \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{X}_0)^T)\mathbf{K}_{f^2,p^2}\mathbf{P}_{R_B}\}^T \\
&\quad \times \{\frac{1}{p}(\mathbf{I}_{f^2} \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{A}_0^{-1}))\mathbf{K}_{f^2,p^2}\mathbf{P}_{R_B}\}^T \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0) \\
&= \text{tr}(\mathbf{B}_0)^2 \frac{\text{tr}(\mathbf{A}_0^2)}{p} \\
T_6 &= \text{vec}(\{\frac{1}{p}(\mathbf{I}_{f^2} \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{X}_0)^T)\mathbf{K}_{f^2,p^2}\mathbf{P}_{R_B}\}^T \\
&\quad \times \{\frac{-1}{pf}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{X}_0)^T \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{Y}_0)^T)\mathbf{K}_{p^2,f^2}\mathbf{P}_{R_A}\}^T \\
&\quad \times \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0) \\
T_7 &= \text{vec}(\{\frac{-1}{pf}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{X}_0)^T \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{Y}_0)^T)\mathbf{K}_{p^2,f^2}\mathbf{P}_{R_A}\}^T \\
&\quad \times \{\frac{1}{f}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{Y}_0)^T \otimes \mathbf{I}_{p^2})\mathbf{P}_{R_A}\}^T \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0) \\
T_8 &= \text{vec}(\{\frac{-1}{pf}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{X}_0)^T \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{Y}_0)^T)\mathbf{K}_{p^2,f^2}\mathbf{P}_{R_A}\}^T \\
&\quad \times \{\frac{1}{p}(\mathbf{I}_{f^2} \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{X}_0)^T)\mathbf{K}_{f^2,p^2}\mathbf{P}_{R_B}\}^T \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0) \\
T_9 &= \text{vec}(\{\frac{-1}{pf}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{X}_0)^T \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{Y}_0)^T)\mathbf{K}_{p^2,f^2}\mathbf{P}_{R_A}\}^T \\
&\quad \times \{\frac{-1}{pf}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{X}_0)^T \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{Y}_0)^T)\mathbf{K}_{p^2,f^2}\mathbf{P}_{R_A}\}^T \\
&\quad \times \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0)
\end{aligned}$$

where we used the definitions of the permutation matrices \mathbf{P}_{R_B} , \mathbf{P}_{R_A} , \mathbf{K}_{p^2,f^2} and \mathbf{K}_{f^2,p^2} .

After some algebra, we observe $T_2 = -T_3 = T_4 = -T_6 = -T_7 = -T_8 = T_9 = \frac{\text{tr}(\mathbf{A}_0^2)\text{tr}(\mathbf{B}_0^2)}{pf}$, which

implies:

$$\text{tr}(\Xi(\Sigma_0 \otimes \Sigma_0)\Xi^T) = T_1 + T_5 + T_7$$

which in turn yields (59).

From (56) and the bounds (57),(58), we have:

$$\begin{aligned} & \left| \frac{n \text{tr}(\text{Cov}(\text{vec}(\tilde{\mathbf{R}}_{FF}(3))))}{\text{tr}(\Xi(\Sigma_0^T \otimes \Sigma_0)\Xi^T)} - 1 \right| \\ & \leq \frac{n(2\|\Xi\|_2 \mathbb{E}[\|\tilde{\mathbf{R}}\|_F \|\Delta_R\|_F] + \mathbb{E}[\|\Delta_R\|_F^2])}{\text{tr}(\Xi(\Sigma_0^T \otimes \Sigma_0)\Xi^T)} \\ & = O\left(\frac{n}{p^2 + f^2} \left\{ (pf)^{3/2} \bar{M} \left(\frac{\log M}{n}\right)^{3/2} \right. \right. \\ & \quad \left. \left. + pf \bar{M}^2 \left(\frac{\log M}{n}\right)^2 \right\} \right) \end{aligned} \quad (62)$$

$$\begin{aligned} & = O\left(\frac{(pf)^{3/2} \bar{M} (\log M)^{3/2}}{p^2 + f^2 \sqrt{n}}\right) \\ & = o(1) \end{aligned} \quad (63)$$

where we used (61) in (62). Note that we used the growth bound assumption $\frac{\bar{M}\sqrt{\log M}}{\sqrt{pf}} = O(\sqrt{n})$ in (63)⁹ and the last step follows from the growth assumption $\frac{(pf)^3 \bar{M}^2 (\log M)^3}{(p^2 + f^2)^2} = o(n)$.

Thus, in first order, we have:

$$\text{tr}(\text{Cov}(\text{vec}(\tilde{\mathbf{R}}_{FF}(3)))) \cong \frac{\text{tr}(\Xi(\Sigma_0 \otimes \Sigma_0)\Xi^T)}{n}$$

as $p, f, n \rightarrow \infty$, which implies:

$$\|\tilde{\mathbf{R}}_{FF}(3)\|_F^2 = O_P\left(\frac{\text{tr}(\Xi(\Sigma_0 \otimes \Sigma_0)\Xi^T)}{n}\right)$$

as $p, f, n \rightarrow \infty$. Combining this with (60), we conclude that the proof for $k = 3$ FF iterations is complete.

Next, we turn to the $k \geq 3$ case, where k is odd. For this case, the main arguments will be sketched since the details follow from similar analysis as before. We saw above that Δ_R is asymptotically negligible as $p, f, n \rightarrow \infty$, so we introduce the approximation symbol “ \cong ” that will retain only the first order terms.

As before, define $\mathbf{B}^0 = \hat{\mathbf{B}}(\mathbf{A}_{init}) - \mathbf{B}_*$ and for $k \geq 3$ odd, recursively define:

$$\text{vec}(\tilde{\mathbf{A}}^{k-2}) \cong \frac{1}{f} \tilde{\mathbf{R}}_A \text{vec}(\mathbf{B}_*^{-1}) - \frac{1}{f} \mathbf{R}_A (\mathbf{B}_*^{-T} \otimes \mathbf{B}_*^{-1}) \text{vec}(\tilde{\mathbf{B}}^{k-3}) \quad (64)$$

$$\text{vec}(\tilde{\mathbf{B}}^{k-1}) \cong \frac{1}{p} \tilde{\mathbf{R}}_B \text{vec}(\mathbf{A}_*^{-1}) - \frac{1}{p} \mathbf{R}_B (\mathbf{A}_*^{-T} \otimes \mathbf{A}_*^{-1}) \text{vec}(\tilde{\mathbf{A}}^{k-2}) \quad (65)$$

⁹This follows since $\frac{\bar{M}\sqrt{\log M}}{\sqrt{pf}} \leq \frac{\bar{M}\sqrt{\log M}}{\sqrt{f/p + \sqrt{p/f}}} = O(\sqrt{n})$.

This is a continuation of the recursive pattern that arises when ignoring the higher-order asymptotics-i.e., for $k = 3$, (64) becomes identical to (38) when Δ_2 is negligible and (65) becomes identical to (46) when Δ_4 is negligible.

Expanding out $\text{vec}(\tilde{\mathbf{A}}^{k-2})$ and $\text{vec}(\tilde{\mathbf{B}}^{k-1})$ for odd $k \geq 3$, we obtain:

$$\begin{aligned} \text{vec}(\tilde{\mathbf{A}}^{k-2}) &\cong \frac{1}{f} \tilde{\mathbf{R}}_A \text{vec}(\mathbf{B}_*^{-1}) \\ &\quad - \left(\frac{k-3}{2}\right) \frac{1}{pf} \text{vec}(\mathbf{A}_*) \text{vec}(\mathbf{B}_*^{-T})^T \tilde{\mathbf{R}}_B \text{vec}(\mathbf{A}_*^{-1}) \\ &\quad + \left(\frac{k-3}{2}\right) \frac{1}{pf} \text{vec}(\mathbf{A}_*) \text{vec}(\mathbf{A}_*^{-T})^T \tilde{\mathbf{R}}_A \text{vec}(\mathbf{B}_*^{-1}) \\ &\quad - \frac{1}{pf} \text{vec}(\mathbf{A}_*) \text{vec}(\mathbf{B}_*^{-T})^T \tilde{\mathbf{R}}_B \text{vec}(\mathbf{A}_{init}^{-1}) \end{aligned} \quad (66)$$

$$\begin{aligned} \text{vec}(\tilde{\mathbf{B}}^{k-1}) &\cong \frac{1}{p} \tilde{\mathbf{R}}_B \text{vec}(\mathbf{A}_*^{-1}) \\ &\quad - \left(\frac{k-1}{2}\right) \frac{1}{pf} \text{vec}(\mathbf{B}_*) \text{vec}(\mathbf{A}_*^{-T})^T \tilde{\mathbf{R}}_A \text{vec}(\mathbf{B}_*^{-1}) \\ &\quad + \left(\frac{k-3}{2}\right) \frac{1}{pf} \text{vec}(\mathbf{B}_*) \text{vec}(\mathbf{B}_*^{-T})^T \tilde{\mathbf{R}}_B \text{vec}(\mathbf{A}_*^{-1}) \\ &\quad + \frac{1}{pf} \text{vec}(\mathbf{B}_*) \text{vec}(\mathbf{B}_*^{-T})^T \tilde{\mathbf{R}}_B \text{vec}(\mathbf{A}_{init}^{-1}) \end{aligned} \quad (67)$$

where we repeatedly used the identities (39) (see (106) in [5]).

Then, at the k th odd FF iteration, letting $\tilde{\mathbf{R}}_{FF}(k) = \hat{\mathbf{R}}_{FF}(k) - \Sigma_0$, we have:

$$\begin{aligned} \text{vec}(\mathcal{R}(\tilde{\mathbf{R}}_{FF}(k))) &\cong \text{vec}(\text{vec}(\tilde{\mathbf{A}}^{k-2}) \text{vec}(\mathbf{B}_*)^T) \\ &\quad + \text{vec}(\text{vec}(\mathbf{A}_*) \text{vec}(\tilde{\mathbf{B}}^{k-1})^T) \end{aligned} \quad (68)$$

where we used Eqs. (22),(23) from [5].

Using (66) and (67) in (68) and simplifying:

$$\text{vec}(\mathcal{R}(\tilde{\mathbf{R}}_{FF}(k))) \cong \mathbf{P}_R^{-1} \Xi \text{vec}(\tilde{\mathbf{R}})$$

where Ξ was defined in (53), $\tilde{\mathbf{R}} = \hat{\mathbf{S}}_n - \Sigma_0$, and \mathbf{P}_R is the permutation matrix defined in (24) from [5]. Note that three cancellations occurred in the last step. Using (24) from [5]:

$$\text{vec}(\tilde{\mathbf{R}}_{FF}(k)) \cong \Xi \text{vec}(\tilde{\mathbf{R}}) \quad (69)$$

Since (69) is identical to (54) when Δ_R is negligible, the same analysis following (54) holds and as a result, the same rate holds for odd $k \geq 3$.

For the $k \geq 3$ even case, the analysis changes slightly, but the same rate holds. To see this, first note that for $k \geq 4$ even,

$$\begin{aligned} \text{vec}(\mathcal{R}(\tilde{\mathbf{R}}_{FF}(k))) &\cong \text{vec}(\text{vec}(\tilde{\mathbf{A}}^{k-1})\text{vec}(\mathbf{B}_*)^T) \\ &\quad + \text{vec}(\text{vec}(\mathbf{A}_*)\text{vec}(\tilde{\mathbf{B}}^{k-2})^T) \end{aligned} \quad (70)$$

Using (66) and (67) in (70), we obtain, after some algebra,

$$\begin{aligned} \text{vec}(\mathcal{R}(\tilde{\mathbf{R}}_{FF}(k))) &\cong \frac{1}{f}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{B}_0^{-1})^T \otimes \mathbf{I}_{p^2})\text{vec}(\tilde{\mathbf{R}}_A) \\ &\quad + \frac{1}{p}(\mathbf{I}_{f^2} \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{A}_0^{-1})^T)\text{vec}(\tilde{\mathbf{R}}_B^T) \\ &\quad - \frac{2}{pf}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{B}_0^{-1})^T \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{A}_0^{-1})^T)\text{vec}(\tilde{\mathbf{R}}_B^T) \\ &= \mathbf{P}_R^{-1}\tilde{\mathbf{\Xi}}\text{vec}(\tilde{\mathbf{R}}) \end{aligned}$$

where

$$\begin{aligned} \tilde{\mathbf{\Xi}} &= \mathbf{P}_R \left[\frac{1}{f}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{Y}_0)^T \otimes \mathbf{I}_{p^2})\mathbf{P}_{R_A} \right. \\ &\quad + \frac{1}{p}(\mathbf{I}_{f^2} \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{X}_0)^T)\mathbf{K}_{f^2,p^2}\mathbf{P}_{R_B} \\ &\quad \left. - \frac{2}{pf}(\text{vec}(\mathbf{B}_0)\text{vec}(\mathbf{Y}_0)^T \otimes \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{X}_0)^T)\mathbf{K}_{f^2,p^2}\mathbf{P}_{R_B} \right] \end{aligned}$$

Thus, $\text{vec}(\tilde{\mathbf{R}}_{FF}(k)) \cong \tilde{\mathbf{\Xi}}\text{vec}(\tilde{\mathbf{R}})$. Proceeding similarly as before, we have:

$$\text{tr}(\tilde{\mathbf{\Xi}}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0)\tilde{\mathbf{\Xi}}^T) = \text{vec}(\tilde{\mathbf{\Xi}}^T\tilde{\mathbf{\Xi}})^T \text{vec}(\boldsymbol{\Sigma}_0 \otimes \boldsymbol{\Sigma}_0) = \sum_{i=1}^9 \tilde{T}_i$$

After some algebra, we notice $\tilde{T}_1 = T_1$, $\tilde{T}_2 = T_2$, $\tilde{T}_5 = T_5$, $\tilde{T}_3 = -2\tilde{T}_2$, $\tilde{T}_2 = \tilde{T}_4$, $\tilde{T}_9 = 4T_2$ and $\tilde{T}_6 = \tilde{T}_7 = \tilde{T}_8 = \tilde{T}_3$. Thus, $\sum_{i=1}^9 \tilde{T}_i = T_1 + T_5 - 2T_2$. From this, it is easy to see that we have the upper and lower bounds:

$$\begin{aligned} \text{tr}(\tilde{\Xi}(\Sigma_0 \otimes \Sigma_0)\tilde{\Xi}^T) &\leq (p^2 + f^2)\|\Sigma_0\|_2^2 \\ \text{tr}(\tilde{\Xi}(\Sigma_0 \otimes \Sigma_0)\tilde{\Xi}^T) &\geq (p^2 + f^2)\lambda_{\min}(\mathbf{A}_0)^2\lambda_{\min}(\mathbf{B}_0)^2 - 2\|\Sigma_0\|_2^2 \end{aligned}$$

The remaining of the proof follows similarly as in the $k \geq 3$ odd case. The proof of the theorem is complete. The same rate holds for the error in the precision matrix (for example, see proof of Thm. 4). ■

APPENDIX L

PROOF OF THEOREM 4

Proof: As in the proof of Thm. 1 in [5], let $\mathbf{B}_* = \frac{\text{tr}(\mathbf{A}_0\mathbf{A}_{init}^{-1})}{p}\mathbf{B}_0$ and $\mathbf{A}_* = (\frac{\text{tr}(\mathbf{A}_0\mathbf{A}_{init}^{-1})}{p})^{-1}\mathbf{A}_0$. Note that Assumption 1 implies that $\|\mathbf{B}_*\|_2 = \Theta(1)$ and $\|\mathbf{A}_*\|_2 = \Theta(1)$ as $p, f \rightarrow \infty$. For conciseness, the statement “with probability $1 - \frac{2}{n^2}$ ” will be abbreviated as “w.h.p.”-i.e., with high probability.

For concreteness, we first present the result for $k = 2$ iterations. Then, we generalize the analysis to all finite flip-flop iterations by induction.

The growth assumptions in the theorem imply

$$\max \left\{ p, f, \frac{f^2}{p}, \left(\frac{\sqrt{pf} + f\sqrt{\frac{f}{p}} + p\sqrt{\frac{p}{f}}}{p + f} \right)^2 \right\} \log M \leq C'n \quad (71)$$

for some constant $C' > 0$ large enough¹⁰. In fact, the growth assumption in the theorem statement can be relaxed to (71).

Define intermediate error matrices:

$$\begin{aligned} \tilde{\mathbf{B}}^0 &= \hat{\mathbf{B}}(\mathbf{A}_{init}) - \mathbf{B}_* \\ \tilde{\mathbf{A}}^1 &= \hat{\mathbf{A}}(\hat{\mathbf{B}}(\mathbf{A}_{init})) - \mathbf{A}_* \end{aligned}$$

Define $\mathbf{Y}_* = \mathbf{B}_*^{-1}$ and $\mathbf{X}_* = \mathbf{A}_*^{-1}$. Also, define:

$$\begin{aligned} \mathbf{Y}_1 &= \hat{\mathbf{B}}(\mathbf{A}_{init})^{-1} \\ \mathbf{X}_2 &= \hat{\mathbf{A}}(\hat{\mathbf{B}}(\mathbf{A}_{init}))^{-1} \end{aligned}$$

¹⁰This constant is independent of p, f, n , but may depend on the constants in Assumption 1.

These inverses exist if $n \geq \max\{\frac{p}{f}, \frac{f}{p}\} + 1$ (see [25]). Define the error $\tilde{\Sigma}_{FF}(k) = \Sigma_{FF}(k) - \Sigma_0$ for $k \geq 2$. For notational simplicity, let $\mathbf{B}_0^{max} := \max_k [\mathbf{B}_0]_{k,k}$ and $\mathbf{A}_0^{max} := \max_i [\mathbf{A}_0]_{i,i}$, $\psi_\tau := \psi(\frac{1}{2+\tau})$, where $\psi(\cdot)$ is defined in Lemma 5.

Lemma 5 implies that for

$$n > \frac{8(2+\tau)^2}{\psi_\tau} \log M \quad (72)$$

then with probability $1 - \frac{2}{n^2}$, we have:

$$\|\tilde{\mathbf{B}}^0\|_F \leq 2\sqrt{2\psi_\tau \mathbf{B}_0^{max}} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2 f \sqrt{\frac{\log M}{np}} \quad (73)$$

As in the proof of Thm. 1 in [5], we vectorize the operations (7) and (8):

$$\begin{aligned} \text{vec}(\hat{\mathbf{A}}(\mathbf{B})) &= \frac{1}{f} \hat{\mathbf{R}}_A \text{vec}(\mathbf{B}^{-1}) \\ \text{vec}(\hat{\mathbf{B}}(\mathbf{A})) &= \frac{1}{p} \hat{\mathbf{R}}_B \text{vec}(\mathbf{A}^{-1}) \end{aligned}$$

where $\hat{\mathbf{R}}_A$ and $\hat{\mathbf{R}}_B$ are permuted versions of the sample covariance matrix [5].

Let $\epsilon' > 1$. Note that from (73), for

$$n \geq (\epsilon' 2\sqrt{2\psi_\tau \mathbf{B}_0^{max}} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2)^2 f^2 p^{-1} \log M \quad (74)$$

with probability $1 - \frac{2}{n^2}$, from Weyl's inequality,

$$\begin{aligned} \lambda_{\min}(\hat{\mathbf{B}}(\mathbf{A}_{init})) &= \lambda_{\min}(\tilde{\mathbf{B}}^0 + \mathbf{B}_*) \geq \lambda_{\min}(\mathbf{B}_*) - \|\tilde{\mathbf{B}}^0\|_2 \\ &\geq \lambda_{\min}(\mathbf{B}_*) - \|\tilde{\mathbf{B}}^0\|_F \geq \left(1 - \frac{1}{\epsilon'}\right) \lambda_{\min}(\mathbf{B}_*) \end{aligned}$$

Thus, w.h.p.,

$$\begin{aligned} \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F &= \|\mathbf{Y}_1(\hat{\mathbf{B}}(\mathbf{A}_{init}) - \mathbf{B}_*)\mathbf{Y}_*\|_F \\ &\leq \|\mathbf{Y}_1\|_2 \|\mathbf{Y}_*\|_2 \|\tilde{\mathbf{B}}^0\|_F = \frac{\|\tilde{\mathbf{B}}^0\|_F}{\lambda_{\min}(\mathbf{B}_*) \lambda_{\min}(\hat{\mathbf{B}}(\mathbf{A}_{init}))} \\ &\leq \left(1 - \frac{1}{\epsilon'}\right)^{-1} \|\mathbf{Y}_*\|_2^2 2\sqrt{2\psi_\tau \mathbf{B}_0^{max}} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2 \\ &\quad \times f p^{-1/2} \sqrt{\frac{\log M}{n}} \end{aligned} \quad (75)$$

Using Lemma 5, for

$$n > \frac{8(2+\tau)^2}{\psi_\tau} \log M \quad (76)$$

then, w.h.p.,

$$\begin{aligned}
\|\tilde{\mathbf{R}}_A\|_2 &= \sup_{\|\mathbf{v}\|_2=1} \|\tilde{\mathbf{R}}_A \mathbf{v}\|_2 \leq p \sup_{\|\mathbf{v}\|_2=1} \|\hat{\mathbf{R}}_A \mathbf{v} - \mathbf{R}_A \mathbf{v}\|_\infty \\
&= pf \sup_{\|\mathbf{v}\|_2=1} \left\| \frac{1}{f} \hat{\mathbf{R}}_A \mathbf{v} - \frac{\langle \text{vec}(\mathbf{B}_0), \mathbf{v} \rangle}{f} \text{vec}(\mathbf{A}_0) \right\|_\infty \\
&\leq 2\sqrt{2\psi_\tau} \mathbf{A}_0^{\max} \|\mathbf{B}_0\|_2 p \sqrt{f} \sqrt{\frac{\log M}{n}}
\end{aligned} \tag{77}$$

Expanding $\tilde{\mathbf{A}}^1$:

$$\begin{aligned}
\text{vec}(\tilde{\mathbf{A}}^1) &= \frac{1}{f} \hat{\mathbf{R}}_A \text{vec}(\mathbf{Y}_1) - \text{vec}(\mathbf{A}_*) \\
&= \frac{\text{tr}(\mathbf{B}_0(\mathbf{Y}_1 - \mathbf{Y}_*))}{f} \text{vec}(\mathbf{A}_0) + \text{vec}(\hat{\mathbf{A}}(\mathbf{B}_*) - \mathbf{A}_*) \\
&\quad + \frac{1}{f} \tilde{\mathbf{R}}_A \text{vec}(\mathbf{Y}_1 - \mathbf{Y}_*)
\end{aligned} \tag{78}$$

where we used $\mathbf{R}_A = \text{vec}(\mathbf{A}_0) \text{vec}(\mathbf{B}_0^T)^T$ (see Eq. (91) from [5]).

Now, using the triangle inequality in (78), the bounds (75) and (77), the Cauchy-Schwarz inequality, we obtain w.h.p. (under conditions (72),(74),(76)), after some algebra:

$$\begin{aligned}
\|\tilde{\mathbf{A}}^1\|_F &\leq \sqrt{\frac{p}{f}} \|\mathbf{A}_0\|_2 \|\mathbf{B}_0\|_2 \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F + p |\hat{\mathbf{A}}(\mathbf{B}_*) - \mathbf{A}_*|_\infty \\
&\quad + \frac{1}{f} \|\tilde{\mathbf{R}}_A\|_2 \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F \\
&\leq \tilde{C}_1 (\sqrt{f} + pf^{-1/2}) \sqrt{\frac{\log M}{n}} + \tilde{C}_2 \sqrt{pf} \frac{\log M}{n}
\end{aligned} \tag{79}$$

where

$$\begin{aligned}
\tilde{C}_1 &:= 2\sqrt{2\psi_\tau} \max \left\{ \kappa(\Sigma_0) \left(1 - \frac{1}{\epsilon'}\right)^{-1} \|\Theta_0\|_2 \|\mathbf{A}_{init}\|_2^2 \right. \\
&\quad \left. \times \mathbf{B}_0^{\max} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2, \mathbf{A}_0^{\max} \left| \frac{p}{\text{tr}(\mathbf{A}_0 \mathbf{A}_{init}^{-1})} \right| \right\} \\
\tilde{C}_2 &:= \left(1 - \frac{1}{\epsilon'}\right)^{-1} 8\psi_\tau \mathbf{A}_0^{\max} \mathbf{B}_0^{\max} \|\mathbf{B}_0\|_2 \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2 \|\Theta_0\|_2^2 \|\mathbf{A}_{init}^{-1}\|_2^2
\end{aligned}$$

Let $c_1 > 0$. For

$$n \geq \left(\frac{\tilde{C}_2}{\tilde{C}_1 c_1} \right)^2 p \log M \tag{80}$$

then, from (79), we have w.h.p.

$$\|\tilde{\mathbf{A}}^1\|_F \leq \tilde{C}_1 (1 + c_1) (\sqrt{f} + pf^{-1/2}) \sqrt{\frac{\log M}{n}} \tag{81}$$

Using the permutation operator $\mathcal{R}(\cdot)$ defined in [5]:

$$\begin{aligned} \text{vec}(\mathcal{R}(\tilde{\Sigma}_{FF}(2))) &= \text{vec}(\text{vec}(\tilde{\mathbf{A}}^1)\text{vec}(\mathbf{B}_*)^T) \\ &\quad + \text{vec}(\text{vec}(\mathbf{A}_*)\text{vec}(\tilde{\mathbf{B}}^0)^T) + \text{vec}(\text{vec}(\tilde{\mathbf{A}}^1)\text{vec}(\tilde{\mathbf{B}}^0)^T) \end{aligned} \quad (82)$$

From (73), (82) and $\text{vec}(\tilde{\Sigma}_{FF}(2)) = \mathbf{P}_R \text{vec}(\mathcal{R}(\tilde{\Sigma}_{FF}(2)))$ [5], under conditions (72),(74),(76) and (80), w.h.p.,

$$\begin{aligned} \|\tilde{\Sigma}_{FF}(2)\|_F &\leq \|\tilde{\mathbf{A}}_1\|_F \|\mathbf{B}_*\|_F + \|\mathbf{A}_*\|_F \|\tilde{\mathbf{B}}^0\|_F \\ &\quad + \|\tilde{\mathbf{A}}^1\|_F \|\tilde{\mathbf{B}}^0\|_F \\ &\leq \tilde{C}_3(p+2f) \sqrt{\frac{\log M}{n}} + \tilde{C}_4(f\sqrt{f/p} + \sqrt{pf}) \frac{\log M}{n} \end{aligned} \quad (83)$$

where

$$\begin{aligned} \tilde{C}_3 &:= \max \left\{ \|\mathbf{B}_0\|_2 \left| \frac{\text{tr}(\mathbf{A}_0 \mathbf{A}_{init}^{-1})}{p} \right| \tilde{C}_1(1+c_1), \right. \\ &\quad \left. 2\sqrt{2\psi_\tau \kappa(\mathbf{A}_0)} \|\mathbf{A}_{init}\|_2 \mathbf{B}_0^{max} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2 \right\} \\ \tilde{C}_4 &:= 2\sqrt{2\psi_\tau} \mathbf{B}_0^{max} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2 \tilde{C}_1(1+c_1) \end{aligned}$$

For

$$n \geq \left(\frac{\tilde{C}_4}{\tilde{C}_3 c_2} \right)^2 \frac{(f\sqrt{f/p} + \sqrt{pf})^2}{(p+2f)^2} \log M \quad (84)$$

then, from (83) w.h.p.,

$$\|\tilde{\Sigma}_{FF}(2)\|_F \leq \tilde{C}_3(1+c_2)(p+2f) \sqrt{\frac{\log M}{n}}$$

The proof for $k = 2$ iterations is complete. Using a simple induction, it follows that the rate (17) holds for all k finite.

Next, we show that the convergence rate in the precision matrix Frobenius error is on the same order as the covariance matrix error. Let $\Theta_{FF}(2) := \Sigma_{FF}(2)^{-1}$. From (81), for

$$n > (\epsilon' \|\mathbf{X}_*\|_2 \tilde{C}_1(1+c_1))^2 (\sqrt{f} + pf^{-1/2})^2 \log M$$

then w.h.p.,

$$\begin{aligned} \|\mathbf{X}_2 - \mathbf{X}_*\|_F &\leq \left(1 - \frac{1}{\epsilon'}\right)^{-1} \|\mathbf{X}_*\|_2^2 \tilde{C}_1(1+c_1) \\ &\quad \times (\sqrt{f} + pf^{-1/2}) \sqrt{\frac{\log M}{n}} \end{aligned} \quad (85)$$

Using (75) and (85), we have w.h.p.,

$$\begin{aligned}
\|\Theta_{FF}(2) - \Theta_0\|_F &\leq \|\mathbf{X}_2 - \mathbf{X}_*\|_F \|\mathbf{Y}_*\|_F \\
&\quad + \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F \|\mathbf{X}_*\|_F + \|\mathbf{X}_2 - \mathbf{X}_*\|_F \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F \\
&\leq \tilde{D}_1(2f + p) \sqrt{\frac{\log M}{n}} + \tilde{D}_2(f \sqrt{\frac{f}{p}} + \sqrt{pf}) \frac{\log M}{n}
\end{aligned} \tag{86}$$

where

$$\begin{aligned}
\tilde{D}_1 &= \|\Theta_0\|_2 \left(1 - \frac{1}{\epsilon'}\right)^{-1} \max \left\{ \kappa(\mathbf{A}_0) \|\mathbf{A}_{init}^{-1}\|_2 \tilde{C}_1(1 + c_1), \right. \\
&\quad \left. 2\sqrt{2\psi_\tau} \|\Theta_0\|_2 \|\mathbf{A}_{init}\|_2 \mathbf{B}_0^{max} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2 \right\} \\
\tilde{D}_2 &= (1 - \frac{1}{\epsilon'})^{-2} 2\sqrt{2\psi_\tau} \|\Theta_0\|_2^2 \mathbf{B}_0^{max} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2 \tilde{C}_1(1 + c_1)
\end{aligned}$$

For

$$n > (\frac{\tilde{D}_2}{\tilde{D}_1 d'})^2 \left(\frac{f \sqrt{f/p} + \sqrt{pf}}{2f + p} \right)^2 \log M$$

the bound (86) becomes w.h.p.,

$$\|\Theta_{FF}(2) - \Theta_0\|_F \leq \tilde{D}_1(1 + d')(2f + p) \sqrt{\frac{\log M}{n}}$$

Thus, the same rate $O_P\left(\sqrt{\frac{(p^2 + f^2) \log M}{n}}\right)$ holds for the precision matrix Frobenius error. ■

APPENDIX M

PROOF OF THEOREM 5

Proof: For ease of notation, let $\lambda_X^{(k)} = \frac{\bar{\lambda}_X^{(k)}}{f}$ and $\lambda_Y^{(k)} = \frac{\bar{\lambda}_Y^{(k)}}{p}$ for all $k \geq 1$. First, we show that the first iteration of the KGL algorithm yields a fast statistical convergence rate of $O_P\left(\sqrt{\frac{(p+f) \log M}{n}}\right)$ by appropriately adjusting the regularization parameters. A simple induction finishes the proof. Adopt the notation from the proof of Thm. 4.

Lemma 5 implies that for

$$n \geq \frac{8(2 + \tau)^2}{\psi_\tau} \log M \tag{87}$$

then w.h.p.,

$$|\tilde{\mathbf{B}}^0|_\infty \leq 2\sqrt{2\psi_\tau} \mathbf{B}_0^{max} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2 \sqrt{\frac{\log M}{np}} \tag{88}$$

where $\tilde{\mathbf{B}}^0 = \hat{\mathbf{B}}(\mathbf{A}_{init}) - \mathbf{B}_*$. From Proposition 1 and (88), we obtain w.h.p.,

$$\begin{aligned} \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F &\leq 2\sqrt{2}(1 + c_y)\sqrt{1 + c_{Y_0}}\|\mathbf{Y}_*\|_2^2 \\ &\times 2\sqrt{2\psi_\tau}\mathbf{B}_0^{max}\|\mathbf{A}_{init}^{-1}\mathbf{A}_0\|_2\sqrt{\frac{f\log M}{np}} \end{aligned} \quad (89)$$

where we also used $s_{Y_0} \leq c_{Y_0}f$ and $\mathbf{Y}_1 := \mathbf{G}(\hat{\mathbf{B}}(\mathbf{A}_{init}), \lambda_Y^{(1)}) = \mathbf{B}_1^{-1}$. Note that $fp^{-1}\log M = o(n)$ was used here.

Let $\hat{\mathbf{A}}^1 := \hat{\mathbf{A}}(\mathbf{B}_1) - \mathbf{A}_*$. Then, we have

$$\begin{aligned} \text{vec}(\hat{\mathbf{A}}^1) &= \frac{1}{f}\hat{\mathbf{R}}_A\text{vec}(\mathbf{Y}_1) - \text{vec}(\mathbf{A}_*) \\ &= \frac{1}{f}\hat{\mathbf{R}}_A\text{vec}(\mathbf{Y}_1 - \mathbf{Y}_*) + \frac{1}{f}\tilde{\mathbf{R}}_A\text{vec}(\mathbf{B}_*^{-1}) \\ &= \frac{1}{f}\mathbf{R}_A\text{vec}(\mathbf{Y}_1 - \mathbf{Y}_*) + \text{vec}(\hat{\mathbf{A}}(\mathbf{B}_*) - \mathbf{A}_*) + \frac{1}{f}\tilde{\mathbf{R}}_A\text{vec}(\mathbf{Y}_1 - \mathbf{Y}_*) \\ &= \frac{\text{tr}(\mathbf{B}_0(\mathbf{Y}_1 - \mathbf{Y}_*))}{f}\text{vec}(\mathbf{A}_0) + \text{vec}(\hat{\mathbf{A}}(\mathbf{B}_*) - \mathbf{A}_*) \\ &\quad + \frac{1}{f}\tilde{\mathbf{R}}_A\text{vec}(\mathbf{Y}_1 - \mathbf{Y}_*) \end{aligned} \quad (90)$$

where we used $\mathbf{R}_A = \text{vec}(\mathbf{A}_0)\text{vec}(\mathbf{B}_0^T)^T$ (see Eq. (91) in [5]). Recall the definition of a mixed norm [24]:

$$\|\mathbf{W}\|_{\alpha,\beta} = \sup_{\|\mathbf{v}\|_\alpha=1} \|\mathbf{W}\mathbf{v}\|_\beta \quad (91)$$

From (89) and (91), we have w.h.p.,

$$\begin{aligned} \frac{1}{f}\|\tilde{\mathbf{R}}_A\text{vec}(\mathbf{Y}_1 - \mathbf{Y}_*)\|_\infty &\leq \frac{\|\tilde{\mathbf{R}}_A\|_{2,\infty}}{f}\|\mathbf{Y}_1 - \mathbf{Y}_*\|_F \\ &\leq 8\psi_\tau\mathbf{A}_0^{max}\mathbf{B}_0^{max}\|\mathbf{B}_0\|_2\|\mathbf{A}_{init}^{-1}\mathbf{A}_0\|_2 \\ &\quad \times 2\sqrt{2}(1 + c_y)\sqrt{1 + c_{Y_0}}\|\mathbf{Y}_*\|_2^2 \frac{1}{\sqrt{p}} \frac{\log M}{n} \end{aligned} \quad (92)$$

where we used the bound on the mixed-norm-i.e., for $n > \frac{8(2+\tau)^2}{\psi_\tau} \log M$, then w.h.p. (from Lemma 5),

$$\begin{aligned} \frac{\|\tilde{\mathbf{R}}_A\|_{2,\infty}}{f} &= \frac{1}{f} \sup_{\|\mathbf{v}\|_2=1} \|\tilde{\mathbf{R}}_A\mathbf{v}\|_\infty \\ &= \sup_{\|\mathbf{v}\|_2=1} \left\| \frac{1}{f}\hat{\mathbf{R}}_A\mathbf{v} - \frac{\langle \text{vec}(\mathbf{B}_0), \mathbf{v} \rangle}{f} \text{vec}(\mathbf{A}_0) \right\|_\infty \\ &\leq 2\sqrt{2\psi_\tau}\mathbf{A}_0^{max}\|\mathbf{B}_0\|_2\sqrt{\frac{\log M}{nf}} \end{aligned}$$

From (90), applying the triangle inequality and using the Cauchy-Schwarz inequality, (89), (92), w.h.p.,

$$\begin{aligned}
|\dot{\mathbf{A}}^1|_\infty &\leq \frac{|\text{tr}(\mathbf{B}_0(\mathbf{Y}_1 - \mathbf{Y}_*))|}{f} |\mathbf{A}_0|_\infty + |\hat{\mathbf{A}}(\mathbf{B}_*) - \mathbf{A}_*|_\infty + \frac{1}{f} \|\tilde{\mathbf{R}}_A \text{vec}(\mathbf{Y}_1 - \mathbf{Y}_*)\|_\infty \\
&\leq \frac{\sqrt{f} \|\mathbf{B}_0\|_2 \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F}{f} |\mathbf{A}_0|_\infty + |\hat{\mathbf{A}}(\mathbf{B}_*) - \mathbf{A}_*|_\infty \\
&\quad + \frac{1}{f} \|\tilde{\mathbf{R}}_A \text{vec}(\mathbf{Y}_1 - \mathbf{Y}_*)\|_\infty \\
&\leq \bar{C}_1 \left(\frac{1}{\sqrt{p}} + \frac{1}{\sqrt{f}} \right) \sqrt{\frac{\log M}{n}} + \bar{C}_2 \frac{1}{\sqrt{p}} \frac{\log M}{n}
\end{aligned} \tag{93}$$

where we used Lemma 5. The constants \bar{C}_1 and \bar{C}_2 are given by:

$$\begin{aligned}
\bar{C}_1 &= \max \left\{ \kappa(\Sigma_0) \|\Theta_0\|_2 \|\mathbf{A}_{init}\|_2^2 2\sqrt{2}(1 + c_y) \sqrt{1 + c_{Y_0}} \right. \\
&\quad \times 2\sqrt{2\psi_\tau} \mathbf{B}_0^{max} \|\mathbf{A}_{init}^{-1} \mathbf{A}_0\|_2, \\
&\quad \left. 2\sqrt{2\psi_\tau} \mathbf{A}_0^{max} \left| \frac{p}{\text{tr}(\mathbf{A}_0 \mathbf{A}_{init}^{-1})} \right| \right\} \\
\bar{C}_2 &= 2\sqrt{2\psi_\tau} \mathbf{A}_0^{max} \|\mathbf{B}_0\|_2 \cdot 2\sqrt{2}(1 + c_y) \sqrt{1 + c_{Y_0}} \|\mathbf{Y}_0\|_2^2 \\
&\quad \times \left| \frac{p}{\text{tr}(\mathbf{A}_0 \mathbf{A}_{init}^{-1})} \right|^2 2\sqrt{2\psi_\tau} \mathbf{B}_0^{max} \|\mathbf{A}_0 \mathbf{A}_{init}^{-1}\|_2
\end{aligned}$$

For

$$n > \left(\frac{\bar{C}_2}{\bar{C}_1 c'} \right)^2 \frac{\log M}{(1 + \sqrt{\frac{p}{f}})^2} \tag{94}$$

the bound in (93) simplifies to

$$|\dot{\mathbf{A}}^1|_\infty \leq \bar{C}_1 (1 + c') \left(\frac{1}{\sqrt{p}} + \frac{1}{\sqrt{f}} \right) \sqrt{\frac{\log M}{n}} \tag{95}$$

From Proposition 1 and (95), we obtain w.h.p.:

$$\begin{aligned}
\|\mathbf{X}_1 - \mathbf{X}_*\|_F &\leq 2\sqrt{2}(1 + c_x) \sqrt{1 + c_{X_0}} \|\mathbf{X}_*\|_2^2 \bar{C}_1 (1 + c') \\
&\quad \times \left(1 + \sqrt{\frac{p}{f}} \right) \sqrt{\frac{\log M}{n}}
\end{aligned} \tag{96}$$

where we used $s_{X_0} \leq c_{X_0} p$ and $\mathbf{X}_1 := \mathbf{G}(\hat{\mathbf{A}}(\mathbf{B}_1), \lambda_X^{(1)})$, $\mathbf{X}_* := \mathbf{A}_*^{-1}$. Note that $(1 + \sqrt{p/f})^2 \log M = o(n)$ was used here.

Finally, using (89) and (96), we obtain w.h.p.:

$$\begin{aligned}
& \|\Theta_{KGL}(2) - \Theta_0\|_F = \|\mathbf{X}_1 \otimes \mathbf{Y}_1 - \mathbf{X}_* \otimes \mathbf{Y}_*\|_F \\
& \leq \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F \sqrt{p} \|\mathbf{X}_*\|_2 + \|\mathbf{X}_1 - \mathbf{X}_*\|_F \sqrt{f} \|\mathbf{Y}_*\|_2 \\
& \quad + \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F \|\mathbf{X}_1 - \mathbf{X}_*\|_F \\
& \leq \bar{C}_3(2\sqrt{f} + \sqrt{p}) \sqrt{\frac{\log M}{n}} + \bar{C}_4(1 + \sqrt{\frac{f}{p}}) \frac{\log M}{n}
\end{aligned} \tag{97}$$

where the constants are given by

$$\begin{aligned}
\bar{C}_3 &= 2\sqrt{2} \|\Theta_0\|_2 \max \left\{ (1 + c_y) \sqrt{1 + c_{Y_0}} \|\mathbf{Y}_0\|_2 \left| \frac{p}{\text{tr}(\mathbf{A}_0 \mathbf{A}_{init}^{-1})} \right|, \right. \\
& \quad \left. (1 + c_x) \sqrt{1 + c_{X_0}} \bar{C}_1(1 + c') \kappa(\mathbf{A}_0) \|\mathbf{A}_{init}^{-1}\|_2 \right\} \\
\bar{C}_4 &= 8 \|\Theta_0\|_2^2 (1 + c_y) \sqrt{1 + c_{Y_0}} (1 + c_x) \sqrt{1 + c_{X_0}} \bar{C}_1(1 + c')
\end{aligned}$$

For

$$n > \left(\frac{\bar{C}_4}{\bar{C}_3 c''} \right)^2 \frac{(1 + \sqrt{f/p})^2}{(2\sqrt{f} + \sqrt{p})^2} \log M \tag{98}$$

the bound (97) further becomes:

$$\|\Theta_{KGL}(2) - \Theta_0\|_F \leq \bar{C}_3(1 + c'')(2\sqrt{f} + \sqrt{p}) \sqrt{\frac{\log M}{n}}$$

Note that $\|\Theta_{KGL}(2) - \Theta_0\|_F^2 = O_P \left(\frac{(p+f+\sqrt{pf}) \log M}{n} \right) = O_P \left(\frac{(p+f) \log M}{n} \right)$ as $p, f, n \rightarrow \infty$. This concludes the first part of the proof. The rest of the proof follows by similar bounding arguments coupled with induction. The rate remains the same as the number of iterations increases, but the constant on front changes.

For

$$n > \left(\frac{\bar{C}_4}{\bar{C}_3 c''} \right)^2 \left(\frac{1 + \sqrt{f/p}}{2\sqrt{f} + \sqrt{p}} \right)^2 \log M$$

the bound (97) further becomes:

$$\|\Theta_{KGL}(2) - \Theta_0\|_F \leq \bar{C}_3(1 + c'')(2\sqrt{f} + \sqrt{p}) \sqrt{\frac{\log M}{n}}$$

Note that $\|\Theta_{KGL}(2) - \Theta_0\|_F^2 = O_P \left(\frac{(p+f+\sqrt{pf}) \log M}{n} \right) = O_P \left(\frac{(p+f) \log M}{n} \right)$ as $p, f, n \rightarrow \infty$. This concludes the first part of the proof. The rest of the proof follows by similar bounding arguments coupled with induction. The rate remains the same as the number of iterations increases, but the constant on front changes.

Next, we show that the convergence rate in the covariance matrix Frobenius error is on the same order as the inverse. From (89), for

$$n > (\epsilon' 2\sqrt{2}(1+c_y)\sqrt{1+c_{Y_0}})^2 \|\mathbf{Y}_*\|_2^2 \kappa(\mathbf{Y}_0)^2 f p^{-1} \log M$$

we have w.h.p. $\lambda_{\min}(\mathbf{Y}_1) \geq \lambda_{\min}(\mathbf{Y}_*) - \|\mathbf{Y}_1 - \mathbf{Y}_*\|_F \geq (1 - \frac{1}{\epsilon'}) \lambda_{\min}(\mathbf{Y}_*)$, which in turn implies w.h.p.,

$$\begin{aligned} \|\mathbf{B}_1 - \mathbf{B}_*\|_F &\leq \left(1 - \frac{1}{\epsilon'}\right)^{-1} 2\sqrt{2}(1+c_y)\sqrt{1+c_{Y_0}}\kappa(\mathbf{Y}_0)^2 \\ &\times \sqrt{\frac{f}{p}} \sqrt{\frac{\log M}{n}} \end{aligned} \quad (99)$$

¹¹ Using a similar argument, from (96), for

$$\begin{aligned} n &> \frac{(\epsilon' 2\sqrt{2}(1+c_x)\sqrt{1+c_{X_0}}\bar{C}_1(1+c'))^2}{\lambda_{\min}(\mathbf{A}_*)^2} \kappa(\mathbf{X}_0)^2 \\ &\times (1 + \sqrt{\frac{p}{f}})^2 \log M \end{aligned}$$

we have w.h.p.,

$$\begin{aligned} \|\mathbf{A}_1 - \mathbf{A}_*\|_F &\leq \left(1 - \frac{1}{\epsilon'}\right)^{-1} 2\sqrt{2}(1+c_x)\sqrt{1+c_{X_0}}\bar{C}_1(1+c') \\ &\times \kappa(\mathbf{X}_0)^2 \left(1 + \sqrt{\frac{p}{f}}\right) \sqrt{\frac{\log M}{n}} \end{aligned} \quad (100)$$

where $\mathbf{A}_1 = \mathbf{X}_1^{-1}$.

Let $\Sigma_{KGL}(2) := \Theta_{KGL}(2)^{-1} = \mathbf{A}_1 \otimes \mathbf{B}_1$. Then, w.h.p.,

$$\begin{aligned} \|\Sigma_{KGL}(2) - \Sigma_0\|_F &\leq \|\mathbf{A}_1 - \mathbf{A}_*\|_F \|\mathbf{B}_*\|_F \\ &+ \|\mathbf{B}_1 - \mathbf{B}_*\|_F \|\mathbf{A}_*\|_F + \|\mathbf{A}_1 - \mathbf{A}_*\|_F \|\mathbf{B}_1 - \mathbf{B}_*\|_F \\ &\leq \bar{D}_1(2\sqrt{f} + \sqrt{p}) \sqrt{\frac{\log M}{n}} + \bar{D}_2(1 + \sqrt{\frac{f}{p}}) \frac{\log M}{n} \end{aligned} \quad (101)$$

where the constants are:

$$\begin{aligned} \bar{D}_1 &= (1 - \frac{1}{\epsilon'})^{-1} 2\sqrt{2} \max \left((1+c_x)\sqrt{1+c_{X_0}}\bar{C}_1(1+c') \right. \\ &\quad \left. \times \kappa(\mathbf{X}_0)^2, (1+c_y)\sqrt{1+c_{Y_0}}\kappa(\mathbf{Y}_0)^2 \right) \\ \bar{D}_2 &= 8(1 - \frac{1}{\epsilon'})^{-2} (1+c_x)\sqrt{1+c_{X_0}}(1+c_y)\sqrt{1+c_{Y_0}} \\ &\quad \times \kappa(\mathbf{X}_0)^2 \kappa(\mathbf{Y}_0)^2 \bar{C}_1(1+c') \end{aligned}$$

¹¹Here, $\mathbf{B}_1 = \mathbf{Y}_1^{-1}$ exists since \mathbf{Y}_1 is positive definite (see (10)).

For

$$n > \left(\frac{\bar{D}_2}{\bar{D}_1 d} \right)^2 \left(\frac{1 + \sqrt{\frac{f}{p}}}{2\sqrt{f} + \sqrt{p}} \right)^2 \log M$$

then (101) implies w.h.p.,

$$\|\Sigma_{KGL}(2) - \Sigma_0\|_F \leq \bar{D}_1(1+d)(2\sqrt{f} + \sqrt{p}) \sqrt{\frac{\log M}{n}}$$

Thus, the same rate $O_P \left(\sqrt{\frac{(p+f) \log M}{n}} \right)$ holds for the error in the covariance matrix. ■

REFERENCES

- [1] G. I. Allen and R. Tibshirani, “Transposable regularized covariance models with an application to missing data imputation,” *The Annals of Applied Statistics*, vol. 4, no. 2, pp. 764–790, 2010.
- [2] M. Yuan and Y. Lin, “Model selection and estimation in the gaussian graphical model,” *Biometrika*, vol. 94, pp. 19–35, 2007.
- [3] O. Banerjee, L. E. Ghaoui, and A. d’Aspremont, “Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data,” *Journal of Machine Learning Research*, vol. 9, pp. 485–516, March 2008.
- [4] P. Dutilleul, “The mle algorithm for the matrix normal distribution,” *J. Statist. Comput. Simul.*, vol. 64, pp. 105–123, 1999.
- [5] K. Werner, M. Jansson, and P. Stoica, “On estimation of covariance matrices with Kronecker product structure,” *IEEE Transactions on Signal Processing*, vol. 56, no. 2, February 2008.
- [6] K. Werner and M. Jansson, “Estimation of kronecker structured channel covariances using training data,” in *Proceedings of EUSIPCO*, 2007.
- [7] A. Dawid, “Some matrix-variate distribution theory: notational considerations and a bayesian application,” *Biometrika*, vol. 68, pp. 265–274, 1981.
- [8] A. K. Gupta and D. K. Nagar, *Matrix Variate Distributions*. Chapman Hill, 1999.
- [9] N. Cressie, *Statistics for Spatial Data*. Wiley, New York, 1993.
- [10] J. Yin and H. Li, “Model selection and estimation in the matrix normal graphical model,” *Journal of Multivariate Analysis*, vol. 107, pp. 119–140, 2012.
- [11] K. Yu, J. Lafferty, S. Zhu, and Y. Gong, “Large-scale collaborative prediction using a nonparametric random effects model,” *ICML*, pp. 1185–1192, 2009.
- [12] E. Bonilla, K. M. Chai, and C. Williams, “Multi-task gaussian process prediction,” *Advances in Neural Information Processing Systems*, pp. 153–160, 2008.
- [13] Y. Zhang and J. Schneider, “Learning multiple tasks with a sparse matrix-normal penalty,” *Advances in Neural Information Processing Systems*, vol. 23, pp. 2550–2558, 2010.
- [14] N. Lu and D. Zimmerman, “On likelihood-based inference for a separable covariance matrix,” Statistics and Actuarial Science Dept., Univ. of Iowa, Iowa City, IA, Tech. Rep., 2004.
- [15] J. Friedman, T. Hastie, and R. Tibshirani, “Sparse inverse covariance estimation with the graphical lasso,” *Biostatistics*, vol. 9, no. 3, pp. 432–441, 2008.

- [16] S. Zhou, J. Lafferty, and L. Wasserman, “Time varying undirected graphs,” *Journal of Machine Learning Research*, vol. 80, pp. 295–319, 2010.
- [17] A. Rothman, P. Bickel, E. Levina, and J. Zhu, “Sparse permutation invariant covariance estimation,” *Electronic Journal of Statistics*, vol. 2, pp. 494–515, 2008.
- [18] P. Ravikumar, M. Wainwright, G. Raskutti, and B. Yu, “High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence,” *Advances in Neural Information Processing Systems*, 2008.
- [19] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [20] C.-J. Hsieh, M. A. Sustik, I. S. Dhillon, and P. Ravikumar, “Sparse inverse covariance matrix estimation using quadratic approximation,” *Advances in Neural Information Processing Systems*, vol. 24, 2011.
- [21] R. G. Bartle and D. R. Sherbert, *Introduction to Real Analysis*. John Wiley & Sons, 2000.
- [22] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*. Springer, 1998.
- [23] L. Isserlis, “On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables,” *Biometrika*, vol. 12, 1918.
- [24] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1990.
- [25] N. Lu and D. Zimmerman, “The likelihood ratio test for a separable covariance matrix,” *Statistics and Probability Letters*, vol. 73, no. 5, pp. 449–457, May 2005.